

Dynamic Sampling from Graphical Models

Weiming Feng*

Nisheeth K. Vishnoi[†]

Yitong Yin^{*‡}

Abstract

In this paper, we study the problem of sampling from a graphical model when the model itself is changing dynamically with time. This problem derives its interest from a variety of inference, learning, and sampling settings in machine learning, computer vision, statistical physics, and theoretical computer science. While the problem of sampling from a static graphical model has received considerable attention, theoretical works for its dynamic variants have been largely lacking. The main contribution of this paper is an algorithm that can sample dynamically from a broad class of graphical models over discrete random variables. Our algorithm is parallel and Las Vegas: it knows when to stop and it outputs samples from the exact distribution. We also provide sufficient conditions under which this algorithm runs in time proportional to the size of the update, on general graphical models as well as well-studied specific spin systems. In particular we obtain, for the Ising model (ferromagnetic or anti-ferromagnetic) and for the hardcore model the first dynamic sampling algorithms that can handle both edge and vertex updates (addition, deletion, change of functions), both efficient within regimes that are close to the respective uniqueness regimes, beyond which, even for the static and approximate sampling, no local algorithms were known or the problem itself is intractable. Our dynamic sampling algorithm relies on a local resampling algorithm and a new “equilibrium” property that is shown to be satisfied by our algorithm at each step, and enables us to prove its correctness. This equilibrium property is robust enough to guarantee the correctness of our algorithm, helps us improve bounds on fast convergence on specific models, and should be of independent interest.

This research is supported by the National Key R&D Program of China 2018YFB1003202 and the National Science Foundation of China under Grant Nos. 61722207 and 61672275.

*Nanjing University. Emails: fengwm@smail.nju.edu.cn, yinyt@nju.edu.cn.

[†]École Polytechnique Fédérale de Lausanne (EPFL). Email: nisheeth.vishnoi@gmail.com.

[‡]Part of the work was done when Yitong Yin was visiting the Bernoulli Center at EPFL.

Contents

1	Introduction	1
2	Preliminaries	2
3	Dynamic Sampling	4
3.1	The problem	4
3.2	The dynamic sampler	4
4	Main Results	6
4.1	The equilibrium property	6
4.2	The correctness of the algorithm	8
4.3	The running time of the algorithm	8
4.4	Dynamic sampling from the spin systems	9
5	Related Work	11
6	Proof of Equilibrium and Correctness	11
6.1	Equilibrium conditions	12
6.2	Equilibrium property of resampling algorithms	14
6.3	A meta algorithm for resampling	18
7	Proof of Fast Convergence	19
8	Applications to Spin Systems	22
8.1	Ising model	22
8.2	Hardcore model	25
9	A Lower Bound for Coloring	28
10	Conclusion and Future Work	29

1 Introduction

Graphical models arise in a variety of disciplines ranging from statistical physics, machine learning, statistics, to theoretical computer science. A graphical model is composed of a variable set V and a constraint set E . We consider graphical models on variables with finite support. In this setting, each variable $v \in V$ is associated to a distribution ϕ_v over the set $[q] = \{1, 2, \dots, q\}$. Each constraint $e \in E$ is a subset of variables and comes with a function $\phi_e : [q]^e \rightarrow [0, 1]$ defined on the variables in e . Together, these induce a probability distribution μ over all possible assignments $\sigma \in [q]^V$ as follows:

$$\mu(\sigma) \propto \prod_{v \in V} \phi_v(\sigma_v) \prod_{e \in E} \phi_e(\sigma_e),$$

where σ_v (respectively σ_e) corresponds to the restriction of σ on v (respectively e). This distribution is often referred to as the Gibbs distribution. Such graphical models can capture probability distributions over exponentially sized domains in a succinct manner. The computational problems that arise from the application of graphical models in practice include sampling from the probability distribution they encode, computing marginals (inference), and learning a graphical model; see the books by [38, 40, 49]. These problems often turn out to be computationally hard in the worst case and there is a wide range of methods geared towards solving these problems approximately: Markov chain Monte Carlo (MCMC) methods [39], correlation decay [50], belief propagation [52], and continuous optimization [46].

We focus on the problem of sampling from a graphical model and, in particular, when the model itself is changing *dynamically* with time. For instance, at each time, one or more of the functions ϕ_v or ϕ_e could change. Formally, the computational question that we study is:

*Can we obtain a sample from an updated graphical model
with a small incremental cost?*

This problem captures various settings in computer vision, statistical physics, and machine learning. In computer vision, discrete-valued graphical models are used to represent images and the problem of denoising an image boils down to sampling from such a graphical model. Thus, algorithms that can sample dynamically with a small incremental cost are useful in denoising videos, which can be thought of as a sequence of closely related images; see [3]. As another instance of this question, consider the setting where one uses an optimization algorithm such as stochastic gradient descent (or expectation maximization), to learn a graphical model [36]. Here, the gradient step updates the parameters of the model locally and samples from the updated distribution are used to compute the new gradient. In theoretical computer science, the result of Jerrum, Sinclair, and Vigoda [34] to compute the permanent of a non-negative matrix can also be viewed in this framework: their algorithm starts from a bipartite graph where it is easy to sample a perfect matching and, in each step the bipartite graph is updated. While the problem of sampling from static graphical models has received considerable attention, theoretical works for its dynamic variant that work with general graphical models have been largely lacking. Indeed, since a local update may potentially change significantly the probability space encoded by the graphical model, in general it is unclear whether there should even exist such an algorithmic machinery that can transform with small cost a sample from the original probability space to a new sample from the updated probability space.

In this paper we show there exists such an algorithmic machinery for dynamic sampling from graphical models. The main contribution of this paper is an algorithm that allows us to sample from a broad class of graphical models dynamically. We allow updates (addition, deletion, and changes to the functions) to the variables and constraints. Given a sample from the current graphical model, our sampling algorithm outputs a sample from the updated graphical model. In addition, the algorithm is parallel and Las Vegas; it knows when to stop and it outputs samples from the *exact* distribution. We also provide sufficient conditions under which the algorithm runs in time proportional to the size of the update. This gives the first dynamic sampling algorithm that can handle both variable and constraint updates (addition, deletion, change of functions) for various graphical models. In particular, for the Ising model with inverse temperature β (ferromagnetic or anti-ferromagnetic) and bounded maximum degree Δ under the condition $e^{-2|\beta|} \geq 1 - \frac{1}{2.222\Delta+1}$, and for the hardcore model with fugacity λ and bounded maximum degree Δ under the condition $\lambda \leq \frac{1}{\sqrt{2\Delta-1}}$, we obtain dynamic sampling algorithms that upon each update of an edge or a vertex can draw a new sample within $O(1)$ incremental cost. Meanwhile, in the “non-uniqueness regimes” for these models where respectively $e^{-2|\beta|} < 1 - \frac{2}{\Delta}$ and $\lambda > \frac{e}{\Delta-2}$, even for static and approximate sampling, either there is no local algorithm or the problem itself is intractable.

Our dynamic sampling algorithm uses the idea of “resampling”: Given a starting sample, once the graphical model changes, the part of the sample that is no longer valid is resampled (potentially multiple times). The idea of resampling was crucial to the Moser-Tardos algorithm for constructing a satisfying solution to the Lovász Local Lemma (LLL) [42, 23, 25, 27] and the algorithms for sampling uniformly distributed satisfying solution to the LLL [22, 51]. Prior to our work, it was unknown whether there is such a local resampling rule that can even generate the *correct distribution* for general graphical models. One of our main conceptual contributions is to come up with an “equilibrium” property and show that our resampling algorithm satisfies this property at each step. Roughly, our equilibrium property asserts: conditioning on any subset of variable to be resampled and their current values, the remaining variables are “consistent” with the Gibbs distribution. This property can easily guarantee the correctness of our sampling algorithm in a dynamic setting. Our techniques should be of independent interest and, in particular, could be useful to extend our results to sampling from other spin systems, graphical models over continuous distributions and/or with global constraints.

Organization of this paper. The preliminaries are given in Section 2. In Section 3, we formally define the dynamic sampling problem and give our main algorithm, the *Dynamic Sampler* (Algorithm 1). The main results are stated in Section 4, followed by the related works discussed in Section 5. The correctness and efficiency of the algorithm are analyzed in Section 6 and Section 7 respectively. And these are applied on specific well-studied graphical models in Section 8. A barrier for sampling proper graph coloring is discussed in Section 9. Finally in Section 10, the conclusion and open problems are given.

2 Preliminaries

Graphical models. A (discrete) *graphical model* is a tuple $\mathcal{I} = (V, E, [q], \Phi)$, where V is a set of n variables, $E \subseteq 2^V$ is a set of m constraints (or factors), and $\Phi = (\phi_a)_{a \in E \cup V}$. Each $v \in V$ corresponds to a variable of domain $[q]$ and is associated with a function $\phi_v : [q] \rightarrow \mathbb{R}_{\geq 0}$. Each constraint $e \in E$ is a set of variables with $|e| > 1$, and is associated with a function $\phi_e : [q]^e \rightarrow \mathbb{R}_{\geq 0}$.

Without loss of generality, we assume that each ϕ_v is normalized as a probability distribution over $[q]$, i.e. $\sum_{x \in [q]} \phi_v(x) = 1$; and each ϕ_e is normalized as $\phi_e : [q]^e \rightarrow [0, 1]$. A constraint e is called a *hard* constraint if ϕ_e is Boolean-valued, and otherwise it is called a *soft* constraint.

Each *configuration* $\sigma \in [q]^V$ assigns every variable one of the q possible values, and is assigned following weight:

$$w(\sigma) \triangleq \prod_{v \in V} \phi_v(\sigma_v) \prod_{e \in E} \phi_e(\sigma_e),$$

where σ_e stands for the restriction of σ on subset $e \subseteq V$.

Gibbs distribution. The *Gibbs distribution* of graphical model \mathcal{I} , denoted as $\mu = \mu_{\mathcal{I}}$, is defined as

$$\mu(\sigma) \triangleq \frac{w(\sigma)}{Z}$$

where

$$Z \triangleq \sum_{\sigma \in [q]^V} w(\sigma)$$

is the *partition function*. We simply write $\mu(\sigma) \propto w(\sigma)$. Note that normalization of the constraints as described above does not change the Gibbs distribution μ .

Dependency graphs. The *dependency graph* of the graphical model \mathcal{I} is a graph with vertex set E , where any two constraints $e, f \in E$ are adjacent in the dependency graph if and only if they share a variable. For any constraint $e \in E$, let

$$\Gamma(e) \triangleq \{f \in E \setminus \{e\} \mid f \cap e \neq \emptyset\}$$

denote the neighborhood of e in the dependency graph.

Variable sets. For a subset $D \subseteq V \cup E$ of variables and constraints, we use

$$\text{vbl}(D) \triangleq (D \cap V) \cup \left(\bigcup_{e \in D \cap E} e \right)$$

to denote the set of variables in D or involved in constraints in D .

Internal and boundary constraints. Given a set of variables $S \subseteq V$, we use

$$E(S) \triangleq \{e \in E \mid e \subseteq S\}$$

to denote the set of *internal* constraints defined on variables in S ;

$$\delta(S) \triangleq \{e \in E \mid e \not\subseteq S \wedge e \cap S \neq \emptyset\}$$

the set of *boundary* constraints; and

$$E^+(S) \triangleq E(S) \cup \delta(S)$$

the set of constraints that use variables in S .

3 Dynamic Sampling

We first describe the dynamic sampling setup that we consider in this paper.

3.1 The problem

We consider dynamical graphical model that are subject to local updates. Let $\mathcal{I} = (V, E, [q], \Phi)$ be the input graphical model. We consider following types of local updates:

- **updates for constraints:** modifying the functions ϕ_e of existing constraints $e \in E$; or adding new constraints (e, ϕ_e) where $e \notin E$;
- **updates for variables:** modifying the functions ϕ_v of variables $v \in V$.

We consider the general case where a sequence of updates may be applied to the graphical model simultaneously. An *update request*, or simply an *update*, is represented as a pair (D, Φ_D) . Here $D \subseteq V \cup 2^V$ contains the variables and constraints to be updated, where each $a \in D$ is either a variable $a \in V$, or an existing constraint $a \in E$, or a new constraint $a \in 2^V \setminus E$ with $|e| > 1$; and $\Phi_D = (\phi_a)_{a \in D}$ specifies the function ϕ_a that we are updating to for all $a \in D$.

Remark 3.1 (Deletion and other updates). The deletion of a constraint $e \in E$ can be realized by updating its function ϕ_e to the constant function with value 1. The addition/deletion of independent variables with no incident constraint is trivial to implement. Therefore, without loss of generality we assume the variable set V remains unchanged.

The problem of dynamic sampling from graphical model is then defined as following:

<u>Dynamic Sampling from Graphical Model</u>	
Input:	a graphical model \mathcal{I} , a sample $\mathbf{X} \sim \mu_{\mathcal{I}}$, and an update (D, Φ_D) that modifies \mathcal{I} to \mathcal{I}' ;
Output:	a sample $\mathbf{X}' \sim \mu_{\mathcal{I}'}$.

We assume that the update (D, Φ_D) is fixed arbitrarily by an offline adversary independently of the sample $\mathbf{X} \sim \mu_{\mathcal{I}}$. A stronger adaptive adversary is discussed later in Remark 4.2.

Remark 3.2 (Assumption about the starting sample). The assumption of having a sample from the current graphical model can be easily achieved initially by starting from an empty graphical model on n variables, whose Gibbs distribution is the product distribution $\bigotimes_{v \in V} \phi_v$, after which the availability of such sample is invariant assuming the correctness of dynamic sampling.

3.2 The dynamic sampler

We give a dynamic sampling algorithm (Algorithm 1) for the above problem. The algorithm proceeds by calling a resampling subroutine (Algorithm 2) to resample the variables in $\text{vbl}(D) = (D \cap V) \cup (\bigcup_{e \in D \cap E} e)$, which is the set of all variables involved in the update (D, Φ_D) .

Algorithm 1: Dynamic Sampler

Input : a graphical model \mathcal{I} and a random sample $\mathbf{X} \sim \mu_{\mathcal{I}}$;
Update: an update (D, Φ_D) which modifies \mathcal{I} to \mathcal{I}' ;
Output: a random sample $\mathbf{X} \sim \mu_{\mathcal{I}'}$;
1 $\mathcal{R} \leftarrow \text{vbl}(D)$;
2 **while** $\mathcal{R} \neq \emptyset$ **do**
3 $(\mathbf{X}, \mathcal{R}) \leftarrow \text{Local-Resample}(\mathcal{I}', \mathbf{X}, \mathcal{R})$;
4 **return** \mathbf{X} ;

Algorithm 2: Local-Resample($\mathcal{I}, \mathbf{X}, \mathcal{R}$)

Input : a graphical model $\mathcal{I} = (V, E, [q], \Phi)$, a configuration $\mathbf{X} \in [q]^V$ and a $\mathcal{R} \subseteq V$;
Output: a new pair $(\mathbf{X}', \mathcal{R}')$ of configuration $\mathbf{X}' \in [q]^V$ and subset $\mathcal{R}' \subseteq V$;
1 for each $e \in E^+(\mathcal{R})$, in parallel, compute $\kappa_e \triangleq \frac{1}{\phi_e(\mathbf{X}_e)} \min_{x \in [q]^e: x_{e \cap \mathcal{R}} = \mathbf{X}_{e \cap \mathcal{R}}} \phi_e(x)$;
2 for each $v \in \mathcal{R}$, in parallel, resample $X_v \in [q]$ ind. from the distribution ϕ_v ;
3 for each $e \in E^+(\mathcal{R})$, in parallel, sample $F_e \in \{0, 1\}$ ind. with $\Pr[F_e = 0] = \kappa_e \cdot \phi_e(\mathbf{X}_e)$;
4 $\mathbf{X}' \leftarrow \mathbf{X}$ and $\mathcal{R}' \leftarrow \bigcup_{e \in E: F_e = 1} e$;
5 **return** $(\mathbf{X}', \mathcal{R}')$.

The local resampling procedure. The resampling subroutine, the Local-Resample (Algorithm 2), is the core of our algorithm. For a fixed a graphical model $\mathcal{I} = (V, E, [q], \Phi)$, the resampling procedure takes as input a pair $(\mathbf{X}, \mathcal{R})$ of configuration $\mathbf{X} \in [q]^V$ and subset $\mathcal{R} \subseteq V$ of variables, where \mathcal{R} represents the current *resample set* that contains the “problematic” variables to be resampled. The resampling procedure transforms this input pair $(\mathbf{X}, \mathcal{R})$ to a random pair $(\mathbf{X}', \mathcal{R}')$ of new configuration $\mathbf{X}' \in [q]^V$ and new resample set $\mathcal{R}' \subseteq V$ by the following simple rules:

1. For each problematic variable $v \in \mathcal{R}$, resample its value $X_v \in [q]$ independently according to the distribution ϕ_v . We denote by \mathbf{X}' the configuration resulting from this resampling.
2. For each constraint e affected by the resampling (because some variables in e are resampled), this constraint e is violated ($F_e = 1$ in algorithm) independently with probability $1 - \kappa_e \cdot \phi(\mathbf{X}'_e)$. The variables involved in the violated constraints form the new resample set \mathcal{R}' .

Here $\kappa_e \in [0, 1]$ is a correcting factor computed from the pre-resample configuration \mathbf{X} as:

$$\kappa_e = \frac{1}{\phi_e(\mathbf{X}_e)} \min_{\substack{x \in [q]^e \\ x_{e \cap \mathcal{R}} = \mathbf{X}_{e \cap \mathcal{R}}}} \phi_e(x) \quad (\text{with convention } 0/0 = 1). \quad (1)$$

where the min gives the minimum value of function ϕ_e estimated from observing \mathbf{X} within \mathcal{R} .

Note that κ_e 's are calculated from the configuration \mathbf{X} before the resampling (thus Line 1 and Line 2 in Algorithm 2 are not interchangeable). For the *internal* constraints $e \in E(\mathcal{R})$, κ_e is always 1 thus has no effect on violating such constraints. It may only bias the probabilities of violating the *boundary* constraints $e \in \delta(\mathcal{R})$ by increasing them. Algorithm 1 repeats the above process until the resample set \mathcal{R} is empty.

While our algorithm is simple, establishing its correctness, i.e., it outputs from the right distribution $\mu_{\mathcal{I}}$, is not. Certain steps in the algorithm, for instance, the definition of κ_e 's, are crucial for this purpose and become more clear from the analysis.

Remark 3.3 (Features of the algorithm). Unlike the MCMC sampling, our algorithm is a Las Vegas sampler that knows when it terminates – this is important in simulations. Also, besides being dynamic, our sampling algorithm is parallelizable, and can be implemented as communication-efficient distributed algorithms in a distributed sampling model considered [9, 13].

Remark 3.4 (Comparison with algorithms for constructing and sampling LLL solution). The famous Moser-Tardos algorithm [42, 23, 25, 27] for constructing LLL solution also relies on local resampling of random variables that violate constraints. It was observed by [24, 22] that the Moser-Tardos algorithm does not generate the correct distribution except for very restricted types of constraints. This was fixed by the partial rejection sampling method [22] for uniform sampling LLL solution (graphical models with hard constraints), by resampling an “unblocking” superset of violating random variables (which in our setting corresponds to the case where $\kappa_e = 1$ for all boundary constraints e). A crucial difference between our algorithm and all these previous resampling-based algorithms, is that **our algorithm uses both the current values of the variables and the values after the resampling in determining whether a constraint is violated**. This seems to be a key to sample correctly from general graphical models.

4 Main Results

4.1 The equilibrium property

The correctness and efficiency of our dynamic sampling algorithm rely on an equilibrium property. In this section we present this equilibrium property that is key to our results. First, we introduce the some preliminaries and explain our “conditional Gibbs property”.

Let $\mathcal{I} = (V, E, [q], \Phi)$ be a graphical model with Gibbs distribution $\mu = \mu_{\mathcal{I}}$. Given any $S \subseteq V$ and a *boundary condition* $\tau \in [q]^{V \setminus S}$ such that $\Pr_{X \sim \mu}[X_{V \setminus S} = \tau] > 0$, let μ_S^τ denote the *marginal distribution* induced by μ over S conditioning on τ , i.e.,

$$\forall \sigma \in [q]^S, \quad \mu_S^\tau(\sigma) \triangleq \Pr_{X \sim \mu}[X_S = \sigma \mid X_{V \setminus S} = \tau].$$

We extend the definition of the marginal distribution μ_S^τ to the boundary conditions τ that may locally violate hard constraints. For any $S \subseteq V$ and $\tau \in [q]^{V \setminus S}$, let $\tilde{\mu}_S^\tau$ be the distribution over all $\sigma \in [q]^S$ such that

$$\tilde{\mu}_S^\tau(\sigma) \propto w_S^\tau(\sigma) \triangleq \prod_{v \in S} \phi_v(\sigma_v) \prod_{\substack{e \in E \\ e \cap S \neq \emptyset}} \phi_e((\sigma \cup \tau)_e), \quad (2)$$

if $\sum_{\sigma \in [q]^S} w_S^\tau(\sigma) > 0$. And if $\sum_{\sigma \in [q]^S} w_S^\tau(\sigma) = 0$, $\tilde{\mu}_S^\tau(\sigma)$ is no longer a well-defined distribution, in which case we assume $\tilde{\mu}_S^\tau(\sigma) = 0$ for all $\sigma \in [q]^S$ as a convention. Clearly, $\mu_S^\tau = \tilde{\mu}_S^\tau$ for any feasible boundary condition $\tau \in [q]^{V \setminus S}$ with $\Pr_{X \sim \mu}[X_{V \setminus S} = \tau] > 0$.

Recall that our sampling algorithm maintains a random pair (X, \mathcal{R}) of a configuration $X \in [q]^V$ and a “resample set” $\mathcal{R} \subseteq V$ of problematic variables. In the following we consider the random pair (X, \mathcal{S}) where $\mathcal{S} \triangleq V \setminus \mathcal{R}$ represents the “sanity set” which contains non-problematic variables.

Definition 4.1 (Conditional Gibbs property). A random pair $(X, \mathcal{S}) \in [q]^V \times 2^V$ is said to be *conditionally Gibbs* with respect to \mathcal{I} if for any $S \subseteq V$ and any $\tau \in [q]^{V \setminus S}$ that $\Pr_{(X, \mathcal{S})}[\mathcal{S} = S \wedge X_{V \setminus S} = \tau] > 0$, it holds that $\tilde{\mu}_S^\tau$ is a well-defined probability distribution over $[q]^S$ and

$$\forall \sigma \in [q]^S, \quad \Pr[X_S = \sigma \mid \mathcal{S} = S \wedge X_{V \setminus S} = \tau] = \tilde{\mu}_S^\tau(\sigma), \quad (3)$$

i.e., the distribution of X_S conditioning on $\mathcal{S} = S$ and $X_{V \setminus S} = \tau$ is the same as the marginal Gibbs distribution $\tilde{\mu}_S^\tau$.

This definition states a key property for the random pair (X, \mathcal{R}) generated by the resampling procedure: conditioning on any possible resample set $\mathcal{R} = R$ and its configuration $X_R = \tau$, the variables in $\mathcal{S} = V \setminus \mathcal{R}$ follow the marginal Gibbs distribution μ over \mathcal{S} with boundary condition τ .

For a fixed \mathcal{I} , each call of **Local-Resample** transforms the current pair (X, \mathcal{R}) to a new pair $(X', \mathcal{R}') \in [q]^V \times 2^V$ as: $(X', \mathcal{R}') \leftarrow \text{Local-Resample}(\mathcal{I}, X, \mathcal{R})$. This naturally defines a Markov chain on space $[q]^V \times 2^V$ over states (X, \mathcal{S}) as follows:

Definition 4.2 (The resampling chain). Let $\mathfrak{M}_{\text{Res}}$ denote the Markov chain on space $[q]^V \times 2^V$ over states (X, \mathcal{S}) defined as follows. Each transition $(X, \mathcal{S}) \rightarrow (X', \mathcal{S}')$ of this chain is as:

$$\begin{aligned} (X', \mathcal{R}') &\leftarrow \text{Local-Resample}(\mathcal{I}, X, V \setminus \mathcal{S}); \\ \mathcal{S}' &\leftarrow V \setminus \mathcal{R}'. \end{aligned}$$

The chain stops when $\mathcal{S} = V$. We call this chain $\mathfrak{M}_{\text{Res}}$ the *resampling chain*.

The reason we define this chain using the “sanity set” \mathcal{S} (which is the complement of the “resample set” \mathcal{R} maintained by the resampling algorithm **Local-Resample**) instead of using \mathcal{R} itself is the following: while the resampling algorithm works by fixing the problematic variables within set \mathcal{R} , the analysis should focus on the distribution of non-problematic variables outside \mathcal{R} .

We use \mathfrak{M} to abstractly denote Markov chains (X, \mathcal{S}) on space $[q]^V \times 2^V$. A crucial property to guarantee such chain \mathfrak{M} always sample from the correct Gibbs distribution $\mu_{\mathcal{I}}$ when stops (i.e. conditioning on $\mathcal{S} = V$, the sample X follows $\mu_{\mathcal{I}}$) is the following equilibrium condition.

Condition 1 (The equilibrium condition). *If a random pair $(X, \mathcal{S}) \in [q]^V \times 2^V$ is conditionally Gibbs with respect to graphical model \mathcal{I} , then after one-step transition of \mathfrak{M} , the new pair (X', \mathcal{S}') is also conditionally Gibbs with respect to \mathcal{I} .*

Note that the goal of resampling is to draw a sample X from the correct distribution at the end when the resample set $\mathcal{R} = \emptyset$ (i.e. $\mathcal{S} = V$). This equilibrium condition promises something much stronger: at any step of resampling, even while the current set of problematic variables \mathcal{R} may not be empty, the sample X is always faithful to the correct distribution over the remaining variables. Similar invariant property has been observed for the partial rejection sampling on hard disk model [19, Lemma 3].

If this equilibrium condition indeed holds for the chain $\mathfrak{M}_{\text{Res}}$ defined above, then it implies the correctness of our dynamic sampling algorithm (Algorithm 1). However, this equilibrium condition can be difficult to verify in general and, in Section 6.1, we present a sufficient condition which gives a refined equilibrium condition that implies Condition 1 and is more explicit to verify.

4.2 The correctness of the algorithm

By verifying the equilibrium condition on the resampling chain $\mathfrak{M}_{\text{Res}}$, we show that our dynamic sampler always outputs from the correct distribution.

Theorem 4.1 (Correctness of the dynamic sampling algorithm). *Assuming the input sample $\mathbf{X} \sim \mu_{\mathcal{I}}$, upon termination, Algorithm 1 returns a perfect sample $\mathbf{X}' \sim \mu_{\mathcal{I}'}$.*

In previous resampling-based algorithms [42, 23, 25, 27, 22], the analysis keeps track of an infinite-size table of random variables for resampling. In contrast, the correctness of our dynamic sampler is due to the above equilibrium condition, which provides a better understanding of why the algorithm always outputs from the correct distribution even in a dynamic setting, and also provides new information for analyzing the running time.

In fact, this theorem along with the equilibrium property, are proved for a general class of resampling-based sampling algorithms (formally stated in Section 6.3) that include our dynamic sampler as a special case.

Remark 4.2 (Stronger adversary). The above theorem holds even when the update (D, Φ_D) is provided by an online adaptive adversary satisfying certain locality property: the update (D, Φ_D) may be correlated with the current sample $\mathbf{X} \sim \mu_{\mathcal{I}}$, but conditioning on any particular (D, Φ_D) and any current assignment $X_{\text{vbl}(D)} = \tau$, the distribution of X_S , where $S \triangleq V \setminus \text{vbl}(D)$, is precisely the marginal Gibbs distribution μ_S^τ induced by $\mu_{\mathcal{I}}$ over S . We call such an adversary a *locally adaptive adversary*, since it covers the natural adaptive adversaries where D is constructed incrementally by observing \mathbf{X} inside $\text{vbl}(D)$. Such adversary is stronger than the offline adversary where (D, Φ_D) is fixed arbitrarily independent of the sample $\mathbf{X} \sim \mu_{\mathcal{I}}$.

4.3 The running time of the algorithm

While our algorithm is always correct, for its running time to be efficient, some conditions on the graphical model must be satisfied. The reason is that sampling from graphical models in general is NP-hard, which is true even for static and approximate sampling [35, 14]. The following theorem gives a sufficient condition that guarantees that our dynamic sampling algorithm is efficient and each update takes time proportional to the size of the update. The time complexity of the algorithm is measured by the total number of individual resamplings of variables $X_v \in [q]$ and indicators $F_e \in \{0, 1\}$ made by the algorithm during its execution.

Theorem 4.3 (Fast convergence of the dynamic sampling algorithm). *Let $\mathcal{I} = (V, E, [q], \Phi)$ be a graphical model and (D, Φ_D) an update to \mathcal{I} . Assume that the followings hold for the updated graphical model $\mathcal{I}' = (V, E', [q], \Phi')$. For every $e \in E'$, we have $\phi'_e : [q]^e \rightarrow [B_e, 1]$ for some $0 < B_e \leq 1$, and there is a constant $0 < \delta < 1$ such that*

$$\forall e \in E', \quad B_e \geq \left(1 - \frac{1 - \delta}{d + 1}\right)^{1/2}, \quad (4)$$

where $d \triangleq \max_{e \in E} |\Gamma(e)|$ is the maximum degree of the dependency graph of \mathcal{I}' . Then Algorithm 1 terminates within $O(\log |D|)$ iterations in expectation. Further, if $d = O(1)$ and $\max_{e \in E} |e| = O(1)$, the total number of resamplings is bounded by $O(|D|)$ in expectation.

For general graphical models, if only knowing the parameters B_e 's and d , one cannot expect much improvement over this sufficient condition. For example, for the Ising model (ferromagnetic or anti-ferromagnetic, with arbitrary local fields) on graphs with maximum degree Δ with inverse temperature β , the bound (4) translates to a condition $B_e = e^{-2|\beta|} > 1 - \frac{1}{4\Delta} + o(\frac{1}{\Delta})$; while the sampling problem (even in a static and approximate setting) for the anti-ferromagnetic Ising model in the “non-uniqueness regime”, where $e^{-2|\beta|} < 1 - \frac{2}{\Delta}$, is NP-hard [15].

While this sufficient condition on general graphical models focuses on graphical models with soft constraints, our dynamic sampling algorithm is not restricted to such settings. For specific graphical models with hard constraints, for instance, the hardcore model, we show a regime of fast convergence that improves the previous recent result in [22].

Remark 4.4 (Implications to static sampling). Theorem 4.1 together with Theorem 4.3 provide linear-time static sampling algorithms for graphical models with $d = O(1)$ and $\max_{e \in E} |e| = O(1)$ and satisfying (4), where the algorithm can draw a perfect sample from the graphical model within $O(n)$ cost in expectation by adding constraints onto independent random variables. In contrast, for MCMC sampling there is a $\Omega(n \log n)$ lower bound on the mixing time of local Markov chains [28].

4.4 Dynamic sampling from the spin systems

On spin systems, e.g. the Ising model or the hardcore model, by exploiting the equilibrium condition (Condition 1) we obtain dynamic samplers with improved convergence conditions.

The Ising/Potts model. In the *Ising model* on graph $G = (V, E)$, each edge $e \in E$ is associated with an *inverse temperature* $\beta_e \in \mathbb{R}$. The Gibbs distribution over all configurations $\sigma \in \{-1, +1\}^V$ is defined such that

$$\mu(\sigma) \propto \prod_{e=(u,v) \in E} \exp(\beta_e \sigma_u \sigma_v).$$

The model is called ferromagnetic if all $\beta_e > 0$, and anti-ferromagnetic if all $\beta_e < 0$.

For the Ising model, our dynamic sampler is efficient in a regime described as follows.

Theorem 4.5. *Assuming $e^{-2|\beta_e|} \geq 1 - \frac{1}{\alpha\Delta+1}$ for all $e \in E$, where $\alpha \approx 2.221\dots$ is the root of $\alpha = 1 + \frac{2}{1+\exp(-1/\alpha)}$, there is a dynamic sampling algorithm for the Ising model on graphs with maximum degree $\Delta = O(1)$, that upon updates of k edges and vertices, returns a perfect sample within $O(\log k)$ rounds and $O(k)$ incremental costs in expectation.*

This gives the first fully dynamic Ising sampler that can handle both edge and vertex updates (addition, deletion, change of functions). The algorithm is parallel and in the above regime terminates within $O(\log n)$ rounds with high probability.

This bound is asymptotically tight. In the “non-uniqueness regime” where $e^{-2|\beta|} < 1 - \frac{2}{\Delta}$, for the anti-ferromagnetic Ising model, even static and approximate sampling is intractable [15]; and for the ferromagnetic Ising model, by an argument as in [9] there cannot exist such local and parallel sampling algorithms (even for static and approximate sampling) due to the reconstructibility of the ferromagnetic Ising model in the non-uniqueness regime on locally tree-like graphs [4, 16].

The Ising model is a major subject for static sampling. Previously, the famous Jerrum-Sinclair chain on even subgraphs [11] gives a poly-time approximate sampler for the ferro-Ising model. The

same can also be obtained by a recent result of Guo and Jerrum [21] for the rapid mixing of the random cluster model, which combined with the *coupling from the past* (CFTP) of Propp and Wilson [44] also gives a poly-time perfect sampler for the ferro-Ising model. These Ising samplers work in the entire ferromagnetic regime, but require global translations of the probability space and has large polynomial running times. For local algorithms, the rapid mixing result of Mossel and Sly [43] for the *Glauber dynamics* (a local Markov chain, also known as *heat bath* or *Gibbs sampler*) gives a $O(n \log n)$ -time static and approximate sampler for the ferro-Ising model in the uniqueness regime where $\beta > 0$ and $e^{-2\beta} > 1 - \frac{2}{\Delta}$. For Las Vegas samplers, local algorithms such as the random recycler of Fill and Huber [11] and the bounding chain of Huber [29] give linear- or near-linear time Las Vegas perfect samplers for the (ferro- or anti-ferro-) Ising model in regimes with the form $e^{-2|\beta|} > 1 - O(\frac{1}{\Delta})$. All these algorithms are non-parallel and none of them can deal with fully dynamic updates of edges and vertices (addition, deletion, change of functions).

The *Potts model* is a generalization of the Ising model to non-Boolean states. Each instance of the Potts model is the same as Ising model. The Gibbs distribution is now defined over all configurations $\sigma \in [q]^V$, where $q \geq 2$ gives the number of spin states, such that

$$\mu(\sigma) \propto \prod_{e=(u,v) \in E} \exp(\beta_e \cdot (2\delta(\sigma_u, \sigma_v) - 1)),$$

where $\delta(\cdot, \cdot)$ is the Kronecker delta.

For the Potts models, the same bounds as in Theorem 4.5 hold.

The hardcore model. In the *hardcore model* on graph $G = (V, E)$, each vertex $v \in V$ is associated with a *fugacity* $\lambda_v > 0$. The Gibbs distribution is defined over all independent sets I of graph G as $\mu(I) \propto \prod_{v \in I} \lambda_v$. We call a $v \in I$ *occupied* and a $v \notin I$ *unoccupied*.

Theorem 4.6. *Assuming $\lambda_v \leq \frac{1}{\sqrt{2\Delta-1}}$ for all $v \in V$, there is a dynamic sampling algorithm for the hardcore model on graphs with maximum degree $\Delta = O(1)$, that upon updates of k edges and vertices, returns a perfect sample within $O(\log k)$ rounds and $O(k)$ incremental costs in expectation.*

This gives the first fully dynamic hardcore sampler that can handle both edge and vertex updates (addition, deletion, change of functions). The algorithm is also parallel and in the above regime terminates within $O(\log n)$ rounds with high probability.

Our bound for the hardcore model is also asymptotically tight. There is a critical threshold $\lambda_c(\Delta) = \frac{(\Delta-1)^{\Delta-1}}{(\Delta-2)^\Delta} \approx \frac{e}{\Delta-2}$ known as the “uniqueness threshold” such that when $\lambda > \lambda_c(\Delta)$ even static and approximate sampling is intractable [15].

Previously, the Glauber dynamics is known to be rapidly mixing for the hardcore model with $\lambda < \lambda_c(\Delta)$ on amenable graphs [17, 50] as well as graphs with large girth and degree [7], and also with $\lambda \leq \frac{2}{\Delta-2}$ [47, 6] on general graphs. And the perfect sampling methods of [11, 29] give $\tilde{O}(n)$ -time perfect samplers also when $\lambda \leq \frac{2}{\Delta-2}$. All these algorithms are non-parallel and none of them can deal with fully dynamic updates of edges and vertices (addition, deletion, change of functions).

To achieve better convergence, our algorithm deviates slightly from Algorithm 1, but still falls into its generalization (formally introduced in Section 6.3). We actually show that a natural dynamic variant of the algorithm in [22] is always correct dynamically and improve their regime of fast convergence from $\lambda \leq \frac{1}{2\sqrt{e\Delta-1}}$ to $\lambda \leq \frac{1}{\sqrt{2\Delta-1}}$.

Coloring. Our algorithm is inefficient on graphical models defined by “truly repulsive” hard constraints, e.g. uniform proper q -coloring. Formally, being “truly repulsive” means that for every constraint $e \in E$, any partial assignment $\sigma_s \in [q]^s$, where $s \subset e$, can be extended to a violating $\sigma_e \in [q]^e$ with $\phi_e(\sigma_e) = 0$. For such graphical models, our algorithm is inefficient because to get a correct sample the algorithm is eventually forced to resample all variables simultaneously. How to overcome this is left as a major open problem.

5 Related Work

The theory of Markov chain Monte Carlo (MCMC) sampling has been extensively studied in computer science, probability theory and statistics (see [39, 32]). There is a substantial body of works on MCMC sampling from various graphical models, e.g. the hardcore model [47, 6, 7], the Ising model [33, 43], and proper q -coloring [31, 45, 48].

Less were known for Las Vegas perfect samplers. Some major results include the coupling from the past (CPTP) method of Propp and Wilson [44], Wilson’s cycle-popping algorithm for uniform spanning tree [51], Fill’s algorithm [10, 12], the random recycler method of Fill and Huber [11], Huber’s bounding chain method [29], and most recently the partial rejection sampling method of Guo, Jerrum, and Liu [22]. See the monograph of Huber for a survey [30].

The idea of resampling was used in the famous Moser-Tardos algorithm for constructing a satisfying solution to the Lovász local lemma (LLL) [42], followed by a line of remarkable works [23, 25, 27]. There is a profound connection between LLL and counting [41, 26]. One would expect to use the idea of resampling for sampling. However, as observed in [24, 22], the Moser-Tardos algorithm does not generate uniformly distributed LLL solutions. This was fixed by the partial rejection sampling method of Guo, Jerrum and Liu [22] which can generate uniformly distributed LLL solutions by resampling a proper “unblocking” superset of violating variables. Retrospectively, Wilson’s cycle-popping algorithm [51, 18] can be interpreted as using this method on spanning trees. Most recently, in a major breakthrough of Guo and Jerrum [20], a long-standing open problem in the area of approximate counting, the network reliability problem, was solved by using this method.

For sampling from a dataset, instead of from an exponential-sized space of configurations, sampling from a dynamically increasing dataset (or data stream) with small maintenance cost is a fundamental problem and is the main purpose of the classical *reservoir sampling* methods [37].

The problem of dynamic sampling from graphical models can also be loosely seen as the sampling variant of the dynamic graph problem. In the dynamic graph problem, edges (constraints) are added or removed over time. The goal is to maintain with small cost for each update, a feasible solution or a locally/globally optimal solution (instead of a random solution as in sampling), to some constraint satisfaction graph problem. The dynamic graph problem has a rich history and is one of the major topics for algorithms and data structures. See [5] for a survey.

6 Proof of Equilibrium and Correctness

Our resampling procedure `Local-Resample` (Algorithm 2) maintains a pair (X, \mathcal{R}) where $X \in [q]^V$ is a configuration and $\mathcal{R} \subseteq V$ is the “resample set” which contains the problematic variables to be resampled. Fixed an instance \mathcal{I} of graphical model, each calling of `Local-Resample` transforms the

current pair (X, \mathcal{R}) to a new pair $(X', \mathcal{R}') \in [q]^V \times 2^V$ as:

$$(X', \mathcal{R}') \leftarrow \text{Local-Resample}(\mathcal{I}, X, \mathcal{R}). \quad (5)$$

Recall the Markov chain $\mathfrak{M}_{\text{Res}}$ on space $[q]^V \times 2^V$, called the *resampling chain*, which is defined in Definition 4.2. The chain $\mathfrak{M}_{\text{Res}}$ is defined over states (X, \mathcal{S}) , where $\mathcal{S} \triangleq V \setminus \mathcal{R}$ stands for the “sanity set” that contains the non-problematic variables. Each transition $(X, \mathcal{S}) \rightarrow (X', \mathcal{S}')$ of chain $\mathfrak{M}_{\text{Res}}$ is derived from the above transition (5) with $\mathcal{S} = V \setminus \mathcal{R}$ and $\mathcal{S}' = V \setminus \mathcal{R}'$.

6.1 Equilibrium conditions

Recall the definition of the *conditional Gibbs property* of a random pair $(X, \mathcal{S}) \in [q]^V \times 2^V$, defined in Definition 4.1, which basically says that conditioning on any fixed set $\mathcal{S} = S$ and any boundary condition $X_{V \setminus S} = \tau$ specified on variables in $V \setminus S$, the distribution of X_S is the same as the marginal Gibbs distribution $\tilde{\mu}_S^\tau$ on S with boundary condition τ (defined as (2)). Note that for the random (X, \mathcal{S}) with this property, in particular, conditioning on $\mathcal{S} = V$ the sample X follows the correct Gibbs distribution $\mu_{\mathcal{I}}$.

We use \mathfrak{M} to abstractly denote Markov chains (X, \mathcal{S}) on space $[q]^V \times 2^V$. Recall the equilibrium condition stated in Condition 1, which says that the property of (X, \mathcal{S}) being conditionally Gibbs with respect to a graphical model \mathcal{I} is invariant under transitions of the chain \mathfrak{M} .

Implication to the correctness of dynamic sampling. Note that if Condition 1 indeed holds for the chain $\mathfrak{M}_{\text{Res}}$, then the correctness of our dynamic sampling algorithm (Algorithm 1) stated in Theorem 4.1 follows directly. This is because for a (X, \mathcal{S}) with the conditional Gibbs property, whenever the algorithm stops (when $\mathcal{S} = V$), the sample X follows the current correct Gibbs distribution. And with Condition 1, the conditional Gibbs property is invariant under resampling Local-Resample.

It only remains to verify that the conditional Gibbs property is also invariant against updates, which is easy by the following argument. Algorithm 1 starts with a sample $X \sim \mu_{\mathcal{I}}$ from the current graphical model \mathcal{I} . And upon update (D, Φ_D) that changes \mathcal{I} to a new instance \mathcal{I}' , in Algorithm 1 we start running the chain $\mathfrak{M}_{\text{Res}}$ with the initial state $(X, V \setminus \text{vbl}(D))$, which is obviously conditionally Gibbs with respect to \mathcal{I} as $X \sim \mu_{\mathcal{I}}$ and $\text{vbl}(D)$ is independent of X . Consequently, this $(X, V \setminus \text{vbl}(D))$ must also be conditionally Gibbs with respect to the new instance \mathcal{I}' , because \mathcal{I} and \mathcal{I}' differ only at functions ϕ_v and ϕ_e for $v, e \in D$, whose definitions do not affect whether $(X, V \setminus \text{vbl}(D))$ is conditionally Gibbs with respect to \mathcal{I} (or \mathcal{I}').

Note that the above argument remains valid even when the update (D, Φ_D) is provided by a locally adaptive adversary as described in Remark 4.2, because with such adversary it still holds that the initial state $(X, V \setminus \text{vbl}(D))$ of the chain $\mathfrak{M}_{\text{Res}}$ is conditionally Gibbs with respect to \mathcal{I} , and the rest follows.

A refined equilibrium condition. In general, the equilibrium condition stated in Condition 1 can still be difficult to verify. Here we give a sufficient condition which implies Condition 1 and is more explicit to verify. Recall that $\tilde{\mu}_S^\tau$ represents the marginal Gibbs distribution on S with boundary condition τ , which is formally defined in (2).

Our refined equilibrium condition is stated as follows.

Condition 2 (Refined equilibrium condition). Let $\mathcal{I} = (V, E, [q], \Phi)$ be a graphical model. Let \mathfrak{M} be a Markov chain on space $[q]^V \times 2^V$ with transition matrix P . For any tuple (S, σ, T, τ) where $S, T \subseteq V$, $\sigma \in [q]^{V \setminus S}$ and $\tau \in [q]^{V \setminus T}$, it holds that

$$\forall \mathbf{y} \in [q]^V \text{ where } y_{V \setminus T} = \tau : \sum_{\substack{\mathbf{x} \in [q]^V \\ x_{V \setminus S} = \sigma}} \tilde{\mu}_S^\sigma(x_S) \cdot P((x, S), (y, T)) = \tilde{\mu}_T^\tau(y_T) \cdot C(S, \sigma, T, \tau),$$

where $C(S, \sigma, T, \tau) \geq 0$ is a finite constant which depends only on (S, σ, T, τ) .

This condition describes a linear system with certain consistency requirement. The equations in the system are grouped according to the tuples (S, σ, T, τ) , where all equations in the same group corresponding to a (S, σ, T, τ) involve only those $x, y \in [q]^V$ with $x_{V \setminus S} = \sigma$ and $y_{V \setminus T} = \tau$ and has the same value on $C(S, \sigma, T, \tau)$. The correct resampling algorithm specified by P is in fact a solution to this system.

Intuitively, this condition guarantees that the equilibrium in Condition 1 holds in the following refined sense. Given a random (X, S) that is conditionally Gibbs, conditioning on any possible $S = S$ and $X_{V \setminus S} = \sigma$, after one-step transition of \mathfrak{M} , the resulting state (Y, T) (starting from the (X, S) generated with the fixed $S = S$ and $X_{V \setminus S} = \sigma$) is still conditionally Gibbs.

It is then easy to verify that Condition 1 is implied by this refined equilibrium condition.

Lemma 6.1 (Sufficiency of Condition 2 to Condition 1). *If a Markov chain \mathfrak{M} on space $[q]^V \times 2^V$ satisfies Condition 2, then it satisfies Condition 1.*

Proof. Let $(\mathbf{X}, \mathcal{S}) \in [q]^V \times 2^V$ be a random pair. Consider a one-step transition of the chain \mathfrak{M} from $(\mathbf{X}, \mathcal{S})$ to $(\mathbf{Y}, \mathcal{T})$. We define

$$H \triangleq \left\{ (S, \sigma) \mid S \subseteq V, \sigma \in [q]^{V \setminus S}, \Pr[X_{V \setminus S} = \sigma \wedge \mathcal{S} = S] > 0 \right\}.$$

Assume that $(\mathbf{X}, \mathcal{S})$ is conditionally Gibbs with respect to \mathcal{I} , which means that for any $(S, \sigma) \in H$, the marginal Gibbs distribution $\tilde{\mu}_S^\sigma$ is a well-defined probability distribution over $[q]^S$ and is also the same as the distribution of X_S conditioning on $\mathcal{S} = S$ and $X_{V \setminus S} = \sigma$. Then for any $T \subseteq V$, $\tau \in [q]^{V \setminus T}$, and any $y \in [q]^V$ that $y_{V \setminus T} = \tau$, we have

$$\begin{aligned} & \Pr[\mathbf{Y} = y \wedge \mathcal{T} = T] \\ &= \sum_{(S, \sigma) \in H} \Pr[X_{V \setminus S} = \sigma \wedge \mathcal{S} = S] \sum_{\substack{\mathbf{x} \in [q]^V \\ x_{V \setminus S} = \sigma}} \tilde{\mu}_S^\sigma(x_S) \cdot P((x, S), (y, T)) \\ &= \tilde{\mu}_T^\tau(y_T) \sum_{(S, \sigma) \in H} C(S, \sigma, T, \tau) \cdot \Pr[X_{V \setminus S} = \sigma \wedge \mathcal{S} = S] \quad (\text{Condition 2}) \\ &= C' \cdot \tilde{\mu}_T^\tau(y_T), \end{aligned}$$

for some $C' = C'(T, \tau)$ which does not depend on y_T . Therefore, fixed any $T \subseteq V$ and any $\tau \in [q]^{V \setminus T}$ which is possible for $(\mathbf{Y}, \mathcal{T})$ (so that $C'(T, \tau) > 0$), the probability $\Pr[Y_T = \cdot \wedge \mathcal{T} = T \wedge Y_{V \setminus T} = \tau]$ is proportional to $\tilde{\mu}_T^\tau(\cdot)$. Formally,

$$\forall \eta \in [q]^T, \quad \Pr[Y_T = \eta \wedge \mathcal{T} = T \wedge Y_{V \setminus T} = \tau] = C' \cdot \tilde{\mu}_T^\tau(\eta). \quad (6)$$

This implies that for any $\eta \in [q]^T$,

$$\begin{aligned}
\Pr[Y_{\mathcal{T}} = \eta \mid \mathcal{T} = T \wedge Y_{V \setminus T} = \tau] &= \frac{\Pr[Y_{\mathcal{T}} = \eta \wedge \mathcal{T} = T \wedge Y_{V \setminus T} = \tau]}{\Pr[\mathcal{T} = T \wedge Y_{V \setminus T} = \tau]} \\
&= \frac{\Pr[Y_{\mathcal{T}} = \eta \wedge \mathcal{T} = T \wedge Y_{V \setminus T} = \tau]}{\sum_{\rho \in [q]^T} \Pr[Y_{\mathcal{T}} = \rho \wedge \mathcal{T} = T \wedge Y_{V \setminus T} = \tau]} \\
&\stackrel{\text{(by (6))}}{=} \frac{C' \cdot \tilde{\mu}_T^\tau(\eta)}{\sum_{\rho \in [q]^T} C' \cdot \tilde{\mu}_T^\tau(\rho)} \\
&\stackrel{\text{(by } C' = C'(T, \tau) > 0)}{=} \tilde{\mu}_T^\tau(\eta).
\end{aligned}$$

This shows that the $Y_{\mathcal{T}}$ conditioning on any possible $\mathcal{T} = T$ and $y_{V \setminus T} = \tau$ follows distribution $\tilde{\mu}_T^\tau$, which means $(\mathbf{Y}, \mathcal{T})$ is conditionally Gibbs with respect to \mathcal{I} . \square

6.2 Equilibrium property of resampling algorithms

We then verify the refined equilibrium condition (Condition 2) on the resampling chain $\mathfrak{M}_{\text{Res}}$. Due to Lemma 6.1, this shows that $\mathfrak{M}_{\text{Res}}$ also satisfies Condition 1, which as discussed, implies the correctness of our dynamic sampler (Algorithm 1), as stated in Theorem 4.1.

Lemma 6.2. *The Markov chain $\mathfrak{M}_{\text{Res}}$ satisfies Condition 2.*

Proof. Let P be the transition matrix of $\mathfrak{M}_{\text{Res}}$. Fix any tuple (S, σ, T, τ) , where $S, T \subseteq V$, $\sigma \in [q]^{V \setminus S}$ and $\tau \in [q]^{V \setminus T}$, and any $y \in [q]^V$ that $y_{V \setminus T} = \tau$. Condition 2 holds if the following always holds

$$\sum_{\substack{x \in [q]^V \\ x_{V \setminus S} = \sigma}} \tilde{\mu}_S^\sigma(x_S) \cdot P((x, S), (y, T)) = \tilde{\mu}_T^\tau(y_T) \cdot C(S, \sigma, T, \tau), \tag{7}$$

where $C(S, \sigma, T, \tau) \geq 0$ depends only on (S, σ, T, τ) .

Observe that Algorithm 2 only resamples the variables in $V \setminus S$, thus we have

$$\forall x \in [q]^V : \quad x_S \neq y_S \implies P((x, S), (y, T)) = 0.$$

Therefore, on the LHS of equation (7), there is only one $x \in [q]^V$ with $x_{V \setminus S} = \sigma$ that may have non-zero $P((x, S), (y, T))$. This $x = x(S, \sigma, y) \in [q]^V$ is uniquely determined as

$$x_v = \begin{cases} \sigma_v & v \in V \setminus S, \\ y_v & v \in S. \end{cases} \tag{8}$$

We simply refer to this $x(S, \sigma, y)$ as x when S, σ and y are clear in the context.

The condition in (7) is then simplified to

$$\tilde{\mu}_S^\sigma(x_S) \cdot P((x, S), (y, T)) = \tilde{\mu}_T^\tau(y_T) \cdot C(S, T, \sigma, \tau), \tag{9}$$

where x is as constructed in (8).

We first calculate the $\tilde{\mu}_S^\sigma(x_S)$ that appears in the LHS of the above equation.

Claim 6.3. *The following holds for the $\tilde{\mu}_S^\sigma(x_S)$ in the LHS of (9):*

$$\tilde{\mu}_S^\sigma(x_S) = C_1 \prod_{v \in S \cap T} \phi_v(y_v) \prod_{e \in \delta(S) \cap E^+(T)} \phi_e(x_e) \prod_{e \in E(S) \cap E^+(T)} \phi_e(y_e),$$

where $C_1 = C_1(S, \sigma, T, \tau) \geq 0$ depends only on S, σ, T, τ .

We then observe that we can make the following assumption without loss of generality:

$$\forall e \in E^+(T) \cap \delta(V \setminus S), \quad \forall z \in [q]^e \text{ that } z_{e \setminus S} = x_{e \setminus S}, \quad \phi_e(z) > 0. \quad (10)$$

If otherwise there exists a $e \in E^+(T) \cap \delta(V \setminus S)$ such that $\phi_e(y_e) = 0$ for some $y_e \in [q]^e$ with $y_{e \setminus S} = x_{e \setminus S}$, then the LHS of equation (9) must have value 0, thus (9) holds trivially with $C(S, T, \sigma, \tau) = 0$. This can be verified by the following two cases. Case.1: $\phi_e(x_e) = 0$, where x is as constructed in (8). Then immediately we have $\tilde{\mu}_S^\sigma(x_S) = 0$. Case.2: $\phi_e(x_e) > 0$. Then in Algorithm 2 with input $(x, V \setminus S)$ we must have $\kappa_e = \min_{z \in [q]^V : z_{e \setminus S} = x_{e \setminus S}} \phi_e(z) = 0$, which means e must be contained in the new resample set, however we have $e \in E^+(T)$, which means $P((x, S), (y, T)) = 0$.

Then, we calculate the transition probability $P((x, S), (y, T))$ under assumption (10).

Claim 6.4. *The following holds for the $P((x, S), (y, T))$ in the LHS of (9):*

$$P((x, S), (y, T)) = C_2 \prod_{v \in T \setminus S} \phi_v(y_v) \prod_{e \in \delta(S) \cap E^+(T)} \frac{\phi_e(y_e)}{\phi_e(x_e)} \prod_{e \in E(V \setminus S) \cap E^+(T)} \phi_e(y_e),$$

where $C_2 = C_2(S, \sigma, T, \tau) \geq 0$ depends only on (S, σ, T, τ) .

The product in above claim is well-defined and has finite value because the ratio $\frac{\phi_e(y_e)}{\phi_e(x_e)}$ has finite value assuming (10).

Combining Claim 6.3 and Claim 6.4, we have

$$\begin{aligned} & \tilde{\mu}_S^\sigma(x_S) \cdot P((x, S), (y, T)) \\ &= C_1 C_2 \prod_{v \in T} \phi_v(y_v) \prod_{e \in E^+(T)} \phi_e(y_e) \\ &= C_1 C_2 \prod_{v \in T} \phi_v(y_v) \prod_{e \in E^+(T)} \phi_e((y_T \cup \tau)_e) \quad (\text{since } y_{V \setminus T} = \tau), \end{aligned}$$

which is precisely $\tilde{\mu}_T^\tau(y_T) \cdot C(S, \sigma, T, \tau)$ for some $C(S, \sigma, T, \tau)$ that depends only on S, σ, T, τ . The first equation is due to $\delta(S), E(S)$ and $E(V \setminus S)$ are disjoint and $\delta(S) \cup E(S) \cup E(V \setminus S) = E$. \square

We then prove Claim 6.3 and Claim 6.4.

Proof of Claim 6.3. Recall that $y_{V \setminus T} = \tau$ and x is constructed as that $x_{V \setminus S} = \sigma$ and $x_S = y_S$. By the definition of the marginal Gibbs distribution, we have

$$\tilde{\mu}_S^\sigma(x_S) = C_{1,1} \prod_{v \in S} \phi_v(x_v) \prod_{e \in E^+(S)} \phi_e(x_e), \quad (11)$$

where $C_{1,1}$ is the reciprocal of the partition function for $\tilde{\mu}_S^\sigma$, which depends only on S, σ .

Note that S can be partitioned into two disjoint subsets $S \setminus T$ and $S \cap T$. Moreover, due to the construction of x and y , we have $x_v = y_v = \tau_v$ for every $v \in S \setminus T$, and $x_v = y_v$ for every $v \in S \cap T$. Therefore,

$$\prod_{v \in S} \phi_v(x_v) = C_{1,2} \prod_{v \in S \cap T} \phi(y_v),$$

where $C_{1,2} = \prod_{v \in S \setminus T} \phi_v(\tau_v)$ depends only on S, T and τ .

Again, note that $E^+(S)$ can be partitioned into two disjoint subsets $E^+(S) \cap E(V \setminus T)$ and $E^+(S) \cap E^+(T)$ because $E(V \setminus T)$ and $E^+(T)$ are complement to each other. Moreover, for every $e \in E(V \setminus T)$, we have x_e fully determined by (S, σ, T, τ) , because $x_{V \setminus S} = \sigma$ and $x_{S \setminus T} = \tau_{S \setminus T}$. Therefore,

$$\prod_{e \in E^+(S)} \phi_e(x_e) = C_{1,3} \prod_{e \in E^+(S) \cap E^+(T)} \phi_e(x_e),$$

where $C_{1,3} = \prod_{e \in E^+(S) \cap E(V \setminus T)} \phi_e(x_e)$ depends only on (S, σ, T, τ) , in particular, $x_v = \sigma_v$ if $v \in V \setminus S$ and $x_v = \tau_v$ if $v \in S \setminus T$.

And the set $E^+(S) \cap E^+(T)$ can be further partitioned into two disjoint subsets $\delta(S) \cap E^+(T)$ and $E(S) \cap E^+(T)$ because $E^+(S) = \delta(S) \uplus E(S)$. Moreover, we have $x_e = y_e$ for every $e \in E(S) \cap E^+(T)$ because $x_S = y_S$. Therefore,

$$\prod_{e \in E^+(S) \cap E^+(T)} \phi_e(x_e) = \prod_{e \in \delta(S) \cap E^+(T)} \phi_e(x_e) \prod_{e \in E(S) \cap E^+(T)} \phi_e(y_e).$$

Combining everything together, we can rewrite (11) as:

$$\tilde{\mu}_S^\sigma(x_S) = C_1 \prod_{v \in S \cap T} \phi(y_v) \prod_{e \in \delta(S) \cap E^+(T)} \phi_e(x_e) \prod_{e \in E(S) \cap E^+(T)} \phi_e(y_e),$$

where $C_1 = C_{1,1} C_{1,2} C_{1,3}$ depends only on (S, σ, T, τ) , which proves the claim. \square

Proof of Claim 6.4. Fix a tuple (S, σ, T, τ) and $y \in [q]^V$ satisfying $y_{V \setminus T} = \tau$. Let $x \in [q]^V$ be constructed as (8). We then calculate the transition probability $P((x, S), (y, T))$.

Recall that the definition of the chain $\mathfrak{M}_{\text{Res}}$ in Definition 4.2. Algorithm 2 takes (x, R) as input where $R \triangleq V \setminus S$. For each $v \in R$ a random value $Y_v \in [q]$ is sampled independently according to the distribution ϕ_v , and for each $v \in S = V \setminus R$, we simply assume $Y_v = x_v$. For each $e \in E^+(R)$, a random value $F_e \in \{0, 1\}$ is sampled independently with $\Pr[F_e = 0] = \kappa_e \phi_e(Y_e)$ where $\kappa_e \triangleq \min_{z \in [q]^e: z_{e \setminus S} = x_{e \setminus S}} \phi_e(z) / \phi_e(x_e)$ (with convention $0/0 = 1$). Finally a random set $\mathcal{R}' \subseteq V$ is constructed as $\mathcal{R}' = \bigcup_{\substack{e \in E^+(R) \\ F_e = 1}} e$.

$P((x, S), (y, T))$ is the probability that $Y = y$ and $\mathcal{R}' = R'$ where $R' \triangleq V \setminus T$, which occurs if and only if following events occur simultaneously.

- $\mathcal{A}_1 : Y_R = y_R;$
- $\mathcal{A}_2 : \exists \mathcal{F} \subseteq E^+(R) \cap E(R'), \text{ s.t. } \text{vbl}(\mathcal{F}) = R' \wedge (\forall e \in \mathcal{F}, F_e = 1);$
- $\mathcal{A}_3 : \forall e \in E^+(R) \setminus E(R'), F_e = 0.$

The first event guarantees that $Y = y$, and the other two events together guarantee that $\mathcal{R}' = R'$. Therefore, by chain rule

$$P((x, S), (y, T)) = \Pr[\mathcal{A}_1 \wedge \mathcal{A}_2 \wedge \mathcal{A}_3] = \Pr[\mathcal{A}_1] \cdot \Pr[\mathcal{A}_2 \mid \mathcal{A}_1] \cdot \Pr[\mathcal{A}_3 \mid \mathcal{A}_1 \wedge \mathcal{A}_2].$$

We then calculate these probabilities separately.

- Since each Y_v for $v \in R$ is sampled independently according to distribution ϕ_v , we have

$$\begin{aligned} \Pr[\mathcal{A}_1] &= \Pr[Y_R = y_R] \\ &= \prod_{v \in R} \phi_v(y_v) \\ &= \prod_{v \in R \cap R'} \phi_v(\tau_v) \prod_{v \in R \setminus R'} \phi_v(y_v) && (y_{R'} = y_{V \setminus T} = \tau) \\ &= C_{2,1} \prod_{v \in T \setminus S} \phi_v(y_v), && (R = V \setminus S \text{ and } R' = V \setminus T) \end{aligned} \quad (12)$$

where $C_{2,1} = C_{2,1}(S, T, \tau) = \prod_{v \in V \setminus (T \cup S)} \phi_v(\tau_v)$ depends only on S, T and τ .

- Event \mathcal{A}_2 can be expressed as the union of a collection of disjoint events, where each disjoint event corresponds to a $\mathcal{F} \subseteq E^+(R) \cap E(R')$ with $\text{vbl}(\mathcal{F}) = R'$ and occurs when \mathcal{F} gives the precise set of e 's with $F_e = 1$. Then, we have

$$\begin{aligned} &\Pr[\mathcal{A}_2 \mid \mathcal{A}_1] \\ &= \sum_{\substack{\mathcal{F} \subseteq E^+(R) \cap E(R') \\ \text{vbl}(\mathcal{F}) = R'}} \Pr[\forall e \in \mathcal{F}, F_e = 1 \wedge \forall e \in (E^+(R) \cap E(R')) \setminus \mathcal{F}, F_e = 0 \mid \mathcal{A}_1] \\ &= \sum_{\substack{\mathcal{F} \subseteq E^+(R) \cap E(R') \\ \text{vbl}(\mathcal{F}) = R'}} \left(\prod_{e \in \mathcal{F}} \Pr[F_e = 1 \mid \mathcal{A}_1] \prod_{\substack{e \in E^+(R) \cap E(R') \\ e \notin \mathcal{F}}} \Pr[F_e = 0 \mid \mathcal{A}_1] \right) \\ &= \sum_{\substack{\mathcal{F} \subseteq E^+(R) \cap E(R') \\ \text{vbl}(\mathcal{F}) = R'}} \left(\prod_{e \in \mathcal{F}} (1 - \kappa_e \cdot \phi_e(\tau_e)) \prod_{\substack{e \in E^+(R) \cap E(R') \\ e \notin \mathcal{F}}} \kappa_e \cdot \phi_e(\tau_e) \right), \end{aligned}$$

where the second equation is due to the conditional independence between F_e 's for all $e \in E^+(R)$ given Y , and the third equation is due to that $y_e = \tau_e$ for all $e \in E(R')$.

We then only need to verify that the κ_e 's in above formula are determined only by (S, σ, T, τ) , which is obvious because x_e for $e \in E(R')$ is determined by these, since $x_R = x_{V \setminus S} = \sigma$ and $x_{R' \setminus R} = x_{S \setminus T} = y_{S \setminus T} = \tau_{S \setminus T}$. Thus, we have

$$\Pr[\mathcal{A}_2 \mid \mathcal{A}_1] = C_{2,2}, \quad (13)$$

for some $C_{2,2} = C_{2,2}(S, \sigma, T, \tau)$ depends only on (S, σ, T, τ) .

- Given Y , all F_e 's are sampled independently. Thus, we have

$$\begin{aligned}
\Pr[\mathcal{A}_3 \mid \mathcal{A}_1 \wedge \mathcal{A}_2] &= \Pr[\mathcal{A}_3 \mid \mathcal{A}_1] \\
&= \prod_{e \in E^+(R) \setminus E(R')} \kappa_e \cdot \phi_e(y_e) \\
&= \prod_{e \in \delta(R) \setminus E(R')} \kappa_e \cdot \phi_e(y_e) \prod_{e \in E(R) \setminus E(R')} \kappa_e \cdot \phi_e(y_e) \\
&= \prod_{e \in \delta(R) \setminus E(R')} \frac{\phi_e(y_e)}{\phi_e(x_e)} \left(\min_{\substack{z \in [q]^e \\ z_e \setminus S = x_e \setminus S}} \phi_e(z) \right) \prod_{e \in E(R) \setminus E(R')} \phi_e(y_e).
\end{aligned}$$

The second equation is due to that $E^+(R) = E(R) \cup \delta(R)$. The last equation is due to the fact that $\kappa_e = 1$ for all $e \in E(R)$. With assumption (10), the ratios in the last equation have finite values.

Note that $R = V \setminus S$ and $R' = V \setminus T$. And also note that for any $e \in \delta(R)$, the value of $\min_{z \in [q]^e: z_e \setminus S = x_e \setminus S} \phi_e(z)$ is determined by S and σ because $x_e \setminus S = \sigma_e \setminus S$. Thus, we have

$$\begin{aligned}
\Pr[\mathcal{A}_3 \mid \mathcal{A}_1 \wedge \mathcal{A}_2] &= C_{2,3} \prod_{e \in \delta(V \setminus S) \setminus E(V \setminus T)} \frac{\phi_e(y_e)}{\phi_e(x_e)} \prod_{e \in E(V \setminus S) \setminus E(V \setminus T)} \phi_e(y_e) \\
&= C_{2,3} \prod_{e \in \delta(S) \cap E^+(T)} \frac{\phi_e(y_e)}{\phi_e(x_e)} \prod_{e \in E(V \setminus S) \cap E^+(T)} \phi_e(y_e) \tag{14}
\end{aligned}$$

where $C_{2,3} = \prod_{e \in \delta(V \setminus S) \setminus E(V \setminus T)} \min_{z \in [q]^e: z_e \setminus S = x_e \setminus S} \phi_e(z)$ depends only on S, σ and T . The second equation is due to $\delta(V \setminus S) = \delta(S)$ and $E(V \setminus T) = E \setminus E^+(T)$.

Combining Equations (12), (13) and (14), we have

$$P((x, S), (y, T)) = C_2 \prod_{v \in T \setminus S} \phi_v(y_v) \prod_{e \in \delta(S) \cap E^+(T)} \frac{\phi_e(y_e)}{\phi_e(x_e)} \prod_{e \in E(V \setminus S) \cap E^+(T)} \phi_e(y_e),$$

where $C_2 = C_{2,1}C_{2,2}C_{2,3}$ depends only on (S, σ, T, τ) . This proves the claim. \square

6.3 A meta algorithm for resampling

Our dynamic sampler can be generalized to a general algorithmic framework for dynamic sampling. In the following we describe a meta-algorithm for resampling called **GenResample**. It takes as input a pair $(\mathbf{X}, \mathcal{R})$ of configuration $\mathbf{X} \in [q]^V$ and a resample set \mathcal{R} , and returns a new random pair $(\mathbf{X}', \mathcal{R}')$. The pseudocode for this meta-algorithm is given in Algorithm 3.

This general resampling procedure consists of two steps: it first expands the current resample set \mathcal{R} to a superset \mathcal{R}'' by calling a subroutine **Expand** $(\mathbf{X}, \mathcal{R})$, which is abstract and may be realized differently in specific implementations; and then the new pair $(\mathbf{X}', \mathcal{R}')$ is constructed by calling **Local-Resample** (Algorithm 2) on the current sample \mathbf{X} and the expanded resample set \mathcal{R}'' .

The *general dynamic sampling algorithm* is obtained by replacing Line 3 of dynamic sampler (Algorithm 1) with $(\mathbf{X}, \mathcal{R}) \leftarrow \text{GenResample}(\mathcal{I}', \mathbf{X}, \mathcal{R})$.

Algorithm 3: GenResample($\mathcal{I}, \mathbf{X}, \mathcal{R}$)

Input : a graphical model $\mathcal{I} = (V, E, [q], \Phi)$, a configuration $\mathbf{X} \in [q]^V$ and a $\mathcal{R} \subseteq V$;
Output: a new pair $(\mathbf{X}', \mathcal{R}')$ of configuration $\mathbf{X}' \in [q]^V$ and subset $\mathcal{R}' \subseteq V$;
1 $\mathcal{R}'' \leftarrow \text{Expand}(\mathbf{X}, \mathcal{R})$;
2 $(\mathbf{X}', \mathcal{R}') \leftarrow \text{Local-Resample}(\mathcal{I}, \mathbf{X}, \mathcal{R}'')$;
3 **return** $(\mathbf{X}', \mathcal{R}')$.

Similar to the definition of $\mathfrak{M}_{\text{Res}}$ corresponding to the Local-Resample algorithm (Definition 4.2), a Markov chain \mathfrak{M}_{GR} on space $[q]^V \times 2^V$ over states (X, \mathcal{S}) , where $\mathcal{S} \triangleq V \setminus \mathcal{R}$ stands for the “sanity set”, can be defined similarly as $\mathfrak{M}_{\text{Res}}$, by replacing Local-Resample with GenResample in Definition 4.2.

We also define another chain $\mathfrak{M}_{\text{Exp}}$ for the Expand subroutine as follows. Each transition $(X, \mathcal{S}) \rightarrow (X', \mathcal{S}')$ is given by:

$$\begin{aligned}\mathcal{R}' &\leftarrow \text{Expand}(X, V \setminus \mathcal{S}); \\ \mathcal{S}' &\leftarrow V \setminus \mathcal{R}'; \\ X' &\leftarrow X.\end{aligned}$$

The \mathfrak{M}_{GR} chain is a composite of the two chains $\mathfrak{M}_{\text{Exp}}$ and $\mathfrak{M}_{\text{Res}}$. Therefore, it is obvious that if $\mathfrak{M}_{\text{Res}}$ and $\mathfrak{M}_{\text{Exp}}$ both satisfy the equilibrium stated in Condition 1, then the composite chain \mathfrak{M}_{GR} also satisfies Condition 1, which as discussed in Section 6.1, would imply the correctness of the general dynamic sampling algorithm with an abstract Expand subroutine. Meanwhile, by Lemma 6.2 and Lemma 6.1, we already has Condition 1 satisfied by the $\mathfrak{M}_{\text{Res}}$ chain. Therefore, the following theorem is true.

Theorem 6.5 (Correctness of the general dynamic sampling algorithm). *Assuming that $\mathfrak{M}_{\text{Exp}}$ satisfies Condition 1 and the input sample $\mathbf{X} \sim \mu_{\mathcal{I}}$, upon termination, the general dynamic sampling algorithm returns a perfect sample $\mathbf{X}' \sim \mu_{\mathcal{I}}$.*

Our dynamic sampler (Algorithm 1) is a special case of the general dynamic sampling algorithm, where the subroutine Expand is the trivial one: $\text{Expand}(\mathbf{X}, \mathcal{R}) = \mathcal{R}$. It is easy to verify that the conditional Gibbs property is invariant under the chain $\mathfrak{M}_{\text{Exp}}$ induced by this trivial Expand subroutine, thus this $\mathfrak{M}_{\text{Exp}}$ satisfies Condition 1. The correctness of Algorithm 1 stated in Thorem 4.1 follows as a corollary of Theorem 6.5.

7 Proof of Fast Convergence

We then analyze the convergence rate of our dynamic sampler on general graphical models.

We first show a contraction behavior of the Local-Resample (Algorithm 2) under condition (4). Fixed a graphical model \mathcal{I} , the resampling procedure Local-Resample takes a pair $(\mathbf{X}, \mathcal{R}) \in [q]^V \times 2^V$ as input and returns a pair $(\mathbf{X}', \mathcal{R}') \in [q]^V \times 2^V$ as

$$(\mathbf{X}', \mathcal{R}') \leftarrow \text{Local-Resample}(\mathbf{X}, \mathcal{R}).$$

We construct an integral-valued *potential function* $H : 2^V \rightarrow \mathbb{Z}_{\geq 0}$ and show that there is a decay on the value of $H(\mathcal{R})$. Let $\mathcal{R} \subseteq V$ be a set of variables. The potential function $H(\mathcal{R})$ is constructed as the minimum number of constraints that can cover all variables in \mathcal{R} . Formally:

Definition 7.1. Let $\mathcal{I} = (V, E, [q], \Phi)$ be a graphical model. The potential function $H : 2^V \rightarrow \mathbb{Z}_{\geq 0}$ is defined as

$$\forall \mathcal{R} \subseteq V : \quad H(\mathcal{R}) \triangleq \min \left\{ |\mathcal{F}| : \mathcal{F} \subseteq E \text{ and } \mathcal{R} \subseteq \bigcup_{e \in \mathcal{F}} e \right\}.$$

Without loss of generality, we assume that each variable is incident to at least one constraint so that the above potential function is well-defined. If otherwise, the variable incident to no constraint is independent with all other variables, which can be easily handled separately by the sampler.

In this definition of potential function, we require that $\mathcal{R} \subseteq \bigcup_{e \in \mathcal{F}} e$ rather than $\mathcal{R} = \bigcup_{e \in \mathcal{F}} e$. This is because in Line 1 of Algorithm 1, the resampling set is constructed as $\mathcal{R} = \text{vbl}(D)$, where $D \subseteq V \cup E$. Since D is allowed to contain some variables in V (for example, D may only contain a single variable), then there may not exist a subset $\mathcal{F} \subseteq E$ such that $\text{vbl}(D) = \bigcup_{e \in \mathcal{F}} e$.

Lemma 7.1. Assume that condition (4) holds for the input graphical model \mathcal{I} . Given any $(\mathbf{X}, \mathcal{R}) \in [q]^V \times 2^V$, the following holds for the $\mathcal{R}' \subseteq V$ returned by Algorithm 2 on the input $(\mathbf{X}, \mathcal{R})$:

$$\mathbb{E} [H(\mathcal{R}')] \leq (1 - \delta)H(\mathcal{R}).$$

Proof. In Algorithm 2, the random configuration X' is obtained by sampling X'_v independently for all $v \in \mathcal{R}$ and setting $X'_v = X_v$ for all $v \notin \mathcal{R}$. Then each constraint $e \in E^+(\mathcal{R})$ resamples $F_e \in \{0, 1\}$ with $\Pr[F_e = 0] = \kappa_e \cdot \phi_e(X'_e)$. Since $\phi_e : [q]^e \rightarrow [B_e, 1]$, then we have

$$\Pr[F_e = 1] \leq 1 - B_e^2. \quad (15)$$

The set \mathcal{R}' returned by Algorithm 2 is constructed by

$$\mathcal{R}' = \bigcup_{e \in E: F_e=1} e = \bigcup_{e \in E^+(\mathcal{R}): F_e=1} e.$$

By the definition of the potential function, $H(\mathcal{R}')$ is at most the number constraint $e \in E^+(\mathcal{R})$ such that $F_e = 1$. By the linearity of expectation and inequality (15), we have

$$\mathbb{E} [H(\mathcal{R}')] \leq \sum_{e \in E^+(\mathcal{R})} \Pr[F_e = 1] \leq \sum_{e \in E^+(\mathcal{R})} (1 - B_e^2). \quad (16)$$

Let $\mathcal{F} \subseteq E$ be the subset of constraints such that $\mathcal{R} \subseteq \bigcup_{e \in \mathcal{F}} e$ and $H(\mathcal{R}) = |\mathcal{F}|$. If the choice of \mathcal{F} is not unique, we pick an arbitrary one. Since $\mathcal{R} \subseteq \bigcup_{e \in \mathcal{F}} e$, then we have $E^+(\mathcal{R}) \subseteq (\bigcup_{e \in \mathcal{F}} \Gamma(e)) \cup \mathcal{F}$. Combining with inequality (16), it holds that

$$\mathbb{E} [H(\mathcal{R}')] \leq \sum_{e \in \mathcal{F}} \left(1 - B_e^2 + \sum_{f \in \Gamma(e)} (1 - B_f^2) \right). \quad (17)$$

By condition (4), it holds that

$$\forall e \in E : \quad B_e^2 \geq 1 - \frac{1 - \delta}{1 + d},$$

where $d = \max_{e \in E} |\Gamma(e)|$ is the maximum degree of dependency graph. Thus

$$\mathbb{E} [H(\mathcal{R}')] \leq \sum_{e \in \mathcal{F}} \left(\frac{1 - \delta}{1 + d} + \sum_{f \in \Gamma(e)} \frac{1 - \delta}{1 + d} \right) \leq (1 - \delta)|\mathcal{F}| = (1 - \delta)H(\mathcal{R}).$$

□

With this stepwise decay on the potentials, we can now prove the main theorem on fast convergence of the dynamic sampler.

Proof of Theorem 4.3. Let $\mathcal{R}_0 = \text{vbl}(D)$ and $\mathbf{X}_0 = \mathbf{X} \sim \mu_{\mathcal{I}}$ be the initial sample. For $t \geq 1$, let

$$(\mathbf{X}_t, \mathcal{R}_t) = \text{Resample}(\mathcal{I}', \mathbf{X}_{t-1}, \mathcal{R}_{t-1}).$$

Let T be the smallest integer such that $\mathcal{R}_T = \emptyset$. The Algorithm 1 terminates after T iterations. By Lemma 7.1, for any $t \geq 1$, we have

$$\mathbb{E} [H(\mathcal{R}_t) \mid \mathcal{R}_{t-1}] \leq (1 - \delta)H(\mathcal{R}_{t-1}).$$

Taking expectation over \mathcal{R}_{t-1} on both sides gives the recurrence:

$$\forall t \geq 1: \quad \mathbb{E} [H(\mathcal{R}_t)] \leq (1 - \delta)\mathbb{E} [H(\mathcal{R}_{t-1})].$$

For the base case, by the definition of potential function, it holds that $H(\mathcal{R}_0) \leq |D|$. With the above recurrence, this implies that for all $t \geq 0$:

$$\mathbb{E} [H(\mathcal{R}_t)] \leq |D|(1 - \delta)^t \leq |D|e^{-t\delta}. \quad (18)$$

Let $\ell = \frac{1}{\delta} \ln |D| + 1$. Then $|D|e^{-\ell\delta} < 1$. By Definition 7.1, it holds that $H(\mathcal{R}) = 0$ if and only if $\mathcal{R} = \emptyset$. We then bound the expected number of iterations as:

$$\begin{aligned} \mathbb{E} [T] &= \sum_{t \geq 1} \Pr[T \geq t] \\ &= \sum_{t \geq 1} \Pr[H(\mathcal{R}_t) \geq 1] \\ &\leq \ell + \sum_{t \geq \ell} |D|e^{-t\delta} \quad (\text{Markov inequality}) \\ &\leq \ell + \frac{1}{1 - e^{-\delta}} \\ &= O(\log |D|). \end{aligned}$$

Hence, the algorithm terminates within $O(\log |D|)$ iterations in expectation.

Recall that $d = \max_{e \in E} |\Gamma(e)|$ is the maximum degree of the dependency graph. Let $k = \max_{e \in E} |e|$ denote the maximum edge size. Then the cost of the t -th iteration in Algorithm 1 is at most $O(kdH(\mathcal{R}_{t-1}))$. This is because the variable subset \mathcal{R}_{t-1} is at most incident to $O(dH(\mathcal{R}_{t-1}))$ constraints. By inequality (18), we have

$$\sum_{t \geq 0} \mathbb{E} [H(\mathcal{R}_t)] \leq \sum_{t \geq 0} |D|e^{-t\delta} = O(|D|).$$

Therefore the expected total cost is bounded by $O(kd|D|)$, which is $O(|D|)$ when the maximum degree d and the maximum constraint size k are both constants. □

8 Applications to Spin Systems

8.1 Ising model

Let $\mathcal{I} = (V, E, \underline{\beta})$ be an Ising model on graph $G = (V, E)$, where each edge $e \in E$ is associated with an *inverse temperature* $\beta_e \in \mathbb{R}$. The Gibbs distribution over all configurations $\sigma \in \{-1, +1\}^V$ is defined such that

$$\mu(\sigma) \propto \prod_{e=(u,v) \in E} \exp(\beta_e \sigma_u \sigma_v).$$

The Ising models can be expressed as graphical models. Our dynamic sampler (Algorithm 1) instantiated on the Ising models gives us the following dynamic perfect Ising sampler (Algorithm 4).

Algorithm 4: Dynamic Ising Sampler

Input : an Ising model \mathcal{I} and a random sample $\mathbf{X} \sim \mu_{\mathcal{I}}$;
Update: an update of vertices and edges in $D \subseteq V \cup \binom{V}{2}$ modifying \mathcal{I} to $\mathcal{I}' = (V, E, \underline{\beta})$;
Output: a random sample $\mathbf{X} \sim \mu_{\mathcal{I}'}$;

- 1 $\mathcal{R} \leftarrow \text{vbl}(D)$;
- 2 **while** $\mathcal{R} \neq \emptyset$ **do**
- 3 all $e \in E$ start with no failure;
- 4 each $e = (u, v) \in \delta(\mathcal{R})$ fails ind. w.p. $1 - e^{-|\beta_e|} \cdot e^{-\beta_e X_u X_v}$;
- 5 for each $v \in \mathcal{R}$, resample $X_v \in \{-1, +1\}$ uniformly and independently;
- 6 each non-failed $e = (u, v) \in E^+(\mathcal{R})$ fails ind. w.p. $1 - e^{-|\beta_e|} \cdot e^{\beta_e X_u X_v}$;
- 7 $\mathcal{R} \leftarrow \bigcup_{\text{failed } e \in E} e$;
- 8 **return** \mathbf{X} ;

For simplicity of exposition we consider the Ising model without external field. For the Ising model with local fields, at Line 5 the random variables are sampled proportional to the local fields, and the algorithm still outputs from the correct Gibbs distribution.

We then prove the fast convergence of this dynamic Ising sampler.

Proof of Theorem 4.5. Let $\mathbf{X} \in \{-1, +1\}^V$ be a random configuration and $\mathcal{R} \subseteq V$ be a random subset of variables such that $(\mathbf{X}, \mathcal{R})$ is conditionally Gibbs with respect to the Ising model \mathcal{I} . Fix one iteration of the while loop in Algorithm 4. The pair $(\mathbf{X}, \mathcal{R})$ is transformed to a random $(\mathbf{X}', \mathcal{R}')$ after this iteration. Assume that $\exp(-2|\beta_e|) \geq 1 - \frac{1}{\alpha\Delta+1}$ for all edge $e \in E$, where $\alpha \approx 2.22\dots$ is the positive root of $\alpha = \frac{2}{1+\exp(-\frac{1}{\alpha})} + 1$. For any fixed \mathcal{R} , we show the following holds

$$\mathbb{E}[H(\mathcal{R}') \mid \mathcal{R}] \leq (1 - \delta(\alpha, \Delta)) H(\mathcal{R}),$$

where the potential function $H(\mathcal{R}) = \min\{|\mathcal{F}| : \mathcal{F} \subseteq E \text{ and } \mathcal{R} \subseteq \bigcup_{e \in \mathcal{F}} e\}$ is as constructed in Definition 7.1 and $\delta(\alpha, \Delta) \triangleq \frac{1}{\alpha(1+\exp(-\frac{1}{\alpha}))(\alpha\Delta+1)}$. Since the above inequality holds for any fixed \mathcal{R} , then we have

$$\mathbb{E}[H(\mathcal{R}')] \leq (1 - \delta(\alpha, \Delta)) \mathbb{E}[H(\mathcal{R})].$$

Note that $\delta(\alpha, \Delta) = \Omega(1)$ since $\Delta = O(1)$. The value of the potential function decays with a constant factor in expectation. Therefore, Theorem 4.5 can be proved by going through the same proof as Theorem 4.1.

The set \mathcal{R}' is constructed by taking the union of all $e \in E^+(\mathcal{R})$ that fail in this iteration. Note that the potential $H(\mathcal{R}')$ is at most the number of failed edges $e \in E^+(\mathcal{R})$. By linearity of expectation,

$$\mathbb{E} [H(\mathcal{R}') \mid \mathcal{R}] \leq \sum_{e \in E^+(\mathcal{R})} \Pr[e \text{ fails}] \quad (19)$$

We then bound the probability that an edge $e \in E^+(\mathcal{R})$ fails. Let β^* be the maximum value of $|\beta_e|$ for all $e \in E$, which is defined as

$$\beta^* \triangleq \max_{e \in E} |\beta_e|.$$

First, consider the internal edges $e = (u, v) \in E(\mathcal{R})$. Note that for such internal edges, $X'_u, X'_v \in \{-1, +1\}$ are sampled uniformly at Line 5 and e only fails at Line 6 with probability $1 - \exp(\beta_e X'_u X'_v - |\beta_e|)$. Therefore,

$$\forall e \in E(\mathcal{R}) : \quad \Pr[e \text{ fails}] = \frac{1 - \exp(-2|\beta_e|)}{2} \leq \frac{1 - \exp(-2\beta^*)}{2}. \quad (20)$$

Consider the boundary edges $e = (u, v) \in \delta(\mathcal{R})$ where $u \in \mathcal{R}$ and $v \notin \mathcal{R}$. To bound the probability that e fails, we give a lower bound of $\Pr[X_v = c]$ for any $c \in \{-1, 1\}$. Let $N(v) \triangleq \{u \in V \mid (u, v) \in E\}$ be the set of neighbors of vertex v in graph G . By the chain rule, we have

$$\forall c \in \{-1, 1\} : \quad \Pr[X_v = c] = \sum_{\sigma \in [q]^{N(v)}} \Pr[X_{N(v)} = \sigma] \Pr[X_v = c \mid X_{N(v)} = \sigma].$$

Note that the pair $(\mathbf{X}, V \setminus \mathcal{R})$ is conditionally Gibbs with respect to \mathcal{I} . Then it holds that

$$\begin{aligned} \forall \sigma \in [q]^{N(v)} : \quad \Pr[X_v = c \mid X_{N(v)} = \sigma] &= \frac{\prod_{u \in N(v)} \exp(\beta_{(u,v)} \sigma_u c)}{\prod_{u \in N(v)} \exp(\beta_{(u,v)} \sigma_u c) + \prod_{u \in N(v)} \exp(-\beta_{(u,v)} \sigma_u c)} \\ &\geq \frac{\prod_{u \in N(v)} \exp(-|\beta_{(u,v)}|)}{\prod_{u \in N(v)} \exp(-|\beta_{(u,v)}|) + \prod_{u \in N(v)} \exp(|\beta_{(u,v)}|)} \\ &\geq \frac{\exp(-\beta^* \Delta)}{\exp(-\beta^* \Delta) + \exp(\beta^* \Delta)}. \end{aligned}$$

Hence, for any edge $e = (u, v) \in \delta(\mathcal{R})$ where $u \in \mathcal{R}$ and $v \notin \mathcal{R}$, we obtain the following bound

$$\forall c \in \{-1, 1\} : \quad \Pr[X_v = c] \geq \frac{\exp(-\beta^* \Delta)}{\exp(-\beta^* \Delta) + \exp(\beta^* \Delta)} = \frac{\exp(-2\beta^* \Delta)}{1 + \exp(-2\beta^* \Delta)}. \quad (21)$$

Note that $X'_v = X_v$ because X_v is not resampled and the edge $e = (u, v) \in \delta(\mathcal{R})$ fails if it fails either at Line 4 or at Line 6. Given \mathbf{X} and \mathbf{X}' the probability that e fails can be expressed as

$$\begin{aligned} \forall e \in \delta(\mathcal{R}) : \quad \Pr[e \text{ fails} \mid \mathbf{X}, \mathbf{X}'] &= 1 - \exp(\beta_e X'_u X'_v - \beta_e X_u X_v - 2|\beta_e|) \\ &= 1 - \exp(\beta_e X_v (X'_u - X_u)) \cdot \exp(-2|\beta_e|). \end{aligned}$$

Since $X'_u \in \{-1, +1\}$ is sampled uniformly and independently at Line 5, we have $\Pr[X'_u = X_u] = \Pr[X'_u \neq X_u] = \frac{1}{2}$. Combining with inequality (21), the probability that e fails can be bounded as

$$\begin{aligned}
\forall e \in \delta(\mathcal{R}) : \quad & \Pr[e \text{ fails}] \\
& \leq 1 - \frac{1}{2} \exp(-2|\beta_e|) - \frac{1}{2} \left(\frac{\exp(-2\beta^* \Delta)}{1 + \exp(-2\beta^* \Delta)} + \frac{\exp(-4|\beta_e|)}{1 + \exp(-2\beta^* \Delta)} \right) \\
& = \frac{1 - \exp(-2|\beta_e|)}{2} + \frac{1 - \exp(-4|\beta_e|)}{2(1 + \exp(-2\beta^* \Delta))} \\
& \leq \frac{1 - \exp(-2\beta^*)}{2} + \frac{1 - \exp(-4\beta^*)}{2(1 + \exp(-2\beta^* \Delta))}. \tag{22}
\end{aligned}$$

Denote $B \triangleq \exp(-2\beta^*)$. Combining inequalities (20) and (22), we have

$$\forall e \in E^+(\mathcal{R}) : \quad \Pr[e \text{ fails}] \leq \frac{1 - B}{2} + \frac{1 - B^2}{2(1 + B^\Delta)} = \frac{(1 - B)}{2} \left(1 + \frac{1 + B}{1 + B^\Delta} \right). \tag{23}$$

Let $\mathcal{F} \subseteq E$ be a subset of edgcs such that $\mathcal{R} \subseteq \bigcup_{e \in \mathcal{F}} e$ and $H(\mathcal{R}) = |\mathcal{F}|$. If such \mathcal{F} is non-unique, we pick an arbitrary one. Since $\mathcal{R} \subseteq \bigcup_{e \in \mathcal{F}} e$, then we have $E^+(\mathcal{R}) \subseteq (\bigcup_{e \in \mathcal{F}} \Gamma(e)) \cup \mathcal{F}$, where $\Gamma(e) = \{f \in E \mid f \neq e \wedge f \cap e \neq \emptyset\}$ denotes the neighborhood of e in the dependency graph. Hence, the size of the set $E^+(\mathcal{R})$ is at most $(2\Delta - 1)|\mathcal{F}|$, because $\max_{e \in E} |\Gamma(e)| = 2(\Delta - 1)$. Combining this fact with (19) and (23), we have

$$\begin{aligned}
\mathbb{E} [H(\mathcal{R}') \mid \mathcal{R}] & \leq \sum_{e \in E^+(\mathcal{R})} \Pr[e \text{ fails}] \\
& \leq \frac{(1 - B)}{2} \left(1 + \frac{1 + B}{1 + B^\Delta} \right) |E^+(\mathcal{R})| \\
& \leq \frac{(2\Delta - 1)(1 - B)}{2} \left(1 + \frac{1 + B}{1 + B^\Delta} \right) |\mathcal{F}|.
\end{aligned}$$

By our assumption, $\exp(-2|\beta_e|) \geq 1 - \frac{1}{\alpha\Delta + 1}$ for all edge $e \in E$, where $\alpha \approx 2.22\dots$ is the positive root of $\alpha = \frac{2}{1 + \exp(-\frac{1}{\alpha})} + 1$, which implies

$$B = \exp(-2\beta^*) = \exp\left(-2 \max_{e \in E} |\beta_e|\right) \geq 1 - \frac{1}{\alpha\Delta + 1}.$$

Note that $|\mathcal{F}| = H(\mathcal{R})$. We have

$$\begin{aligned}
\mathbb{E} [H(\mathcal{R}') \mid \mathcal{R}] & \leq \frac{(2\Delta - 1)(1 - B)}{2} \left(1 + \frac{1 + B}{1 + B^\Delta} \right) H(\mathcal{R}) \\
& \leq \frac{2\Delta - 1}{2(\alpha\Delta + 1)} \left(1 + \frac{2 - \frac{1}{\alpha\Delta + 1}}{1 + \left(1 - \frac{1}{\alpha\Delta + 1}\right)^\Delta} \right) H(\mathcal{R}) \\
& \leq \frac{1}{\alpha} \left(1 + \frac{2 - \frac{1}{\alpha\Delta + 1}}{1 + \exp\left(-\frac{1}{\alpha}\right)} \right) H(\mathcal{R}),
\end{aligned}$$

where the last inequality is due to that $\frac{2\Delta-1}{2(\alpha\Delta+1)} \leq \frac{1}{\alpha}$ and $\left(1 - \frac{1}{\alpha\Delta+1}\right)^\Delta \geq \exp\left(-\frac{1}{\alpha}\right)$, which can be verified to further equal to

$$\begin{aligned} & \left(\frac{1}{\alpha} \left(1 + \frac{2}{1 + \exp\left(-\frac{1}{\alpha}\right)} \right) - \frac{1}{\alpha \left(1 + \exp\left(-\frac{1}{\alpha}\right) \right) (\alpha\Delta + 1)} \right) H(\mathcal{R}) \\ &= \left(1 - \frac{1}{\alpha \left(1 + \exp\left(-\frac{1}{\alpha}\right) \right) (\alpha\Delta + 1)} \right) H(\mathcal{R}), \end{aligned}$$

due to the definition of α .

In conclusion, we prove the desired decay on the potential:

$$\mathbb{E} [H(\mathcal{R}') \mid \mathcal{R}] \leq \left(1 - \frac{1}{\alpha \left(1 + \exp\left(-\frac{1}{\alpha}\right) \right) (\alpha\Delta + 1)} \right) H(\mathcal{R}) = (1 - \delta(\alpha, \Delta)) H(\mathcal{R}).$$

The rest of the proof can be done by going through the same proof as Theorem 4.1. \square

For the Potts model, the Gibbs distribution is defined over all configurations $\sigma \in [q]^V$ such that

$$\mu(\sigma) \propto \prod_{e=(u,v) \in E} \exp(\beta_e \cdot (2\delta(\sigma_u, \sigma_v) - 1)),$$

where $\delta(\cdot, \cdot)$ is the Kronecker delta.

The dynamic Ising sampler (Algorithm 4) can be naturally generalized to the dynamic Potts sampler: at Line 5, each X_v is now resampled from $[q]$ uniformly (or proportional to local fields if there are non-zero fields) and independently, and the failure probabilities at Line 4 and Line 6 are changed to $\exp(-\beta_e \cdot (2\delta(\sigma_u, \sigma_v) - 1) - |\beta_e|)$ and $\exp(\beta_e \cdot (2\delta(\sigma_u, \sigma_v) - 1) - |\beta_e|)$ respectively. It is easy to verify that this algorithm is precisely the dynamic sampler (Algorithm 1) instantiated on the Potts model and the same bounds as in Theorem 4.5 hold for this dynamic Potts sampler.

8.2 Hardcore model

Let $\mathcal{I} = (V, E, \underline{\lambda})$ be a hardcore model on graph $G = (V, E)$, where each vertex $v \in V$ is associated with a fugacity $\lambda_v > 0$. The Gibbs distribution over all configurations $\sigma \in \{0, 1\}^V$ is defined as

$$\mu(\sigma) \propto \begin{cases} \prod_{v \in I(\sigma)} \lambda_v & \text{if } I(\sigma) \text{ is an independent set} \\ 0 & \text{if } I(\sigma) \text{ is not an independent set,} \end{cases}$$

where $I(\sigma) \triangleq \{v \in V \mid \sigma_v = 1\}$.

The hardcore models can be expressed as graphical models. We consider the meta-algorithm for resampling GenResample (Algorithm 3), with the following Expand subroutine:

$$\text{Expand}(\mathbf{X}, \mathcal{R}) \triangleq \mathcal{R} \cup \{v \in V \setminus \mathcal{R} \mid \exists u \in \mathcal{R} \text{ s.t. } (u, v) \in E \wedge X_u = 1\}. \quad (24)$$

On the hardcore model this is instantiated as the dynamic perfect hardcore sampler (Algorithm 5).

Algorithm 5: Dynamic Hardcore Sampler

Input : a hardcore model \mathcal{I} and a random sample $\mathbf{X} \sim \mu_{\mathcal{I}}$;
Update: an update of vertices and edges in $D \subseteq V \cup \binom{V}{2}$ modifying \mathcal{I} to $\mathcal{I}' = (V, E, \underline{\lambda})$;
Output: a random sample $\mathbf{X} \sim \mu_{\mathcal{I}'}$;
1 $\mathcal{R} \leftarrow \text{vbl}(D)$;
2 **while** $\mathcal{R} \neq \emptyset$ **do**
3 for each $v \in \mathcal{R}$ with $X_v = 1$, add all neighbors of v in graph G into \mathcal{R} ;
4 for each $v \in \mathcal{R}$, resample $X_v \in \{0, 1\}$ ind. with $\Pr[X_v = 1] = \frac{\lambda_v}{1 + \lambda_v}$;
5 $\mathcal{R} \leftarrow \bigcup_{\substack{e=(u,v) \in E \\ X_u = X_v = 1}} e$;
6 **return** \mathbf{X} ;

The subroutine $\text{Expand}(\mathbf{X}, \mathcal{R})$ specified in (24) can be implemented as following. Each vertex $v \in \mathcal{R}$ with $X_v = 1$, in parallel, adds all of its neighbors in graph G into the set \mathcal{R} to obtained the expanded resample set $\mathcal{R}'' = \text{Expand}(\mathbf{X}, \mathcal{R})$.

It is easy to verify that the corresponding Markov chain $\mathfrak{M}_{\text{Exp}}$ satisfies the equilibrium condition. The resample subset $\mathcal{R}'' = \text{Expand}(\mathbf{X}, \mathcal{R})$ is fully determined by \mathcal{R} and $X_{\mathcal{R}}$. The subset \mathcal{R}'' gives no extra information about $X_{V \setminus \mathcal{R}''}$. Note that $\mathcal{R} \subseteq \mathcal{R}''$, then $(V \setminus \mathcal{R}'') \subseteq (V \setminus \mathcal{R})$. Therefore, if the pair $(\mathbf{X}, V \setminus \mathcal{R})$ is conditionally Gibbs, then the pair $(\mathbf{X}, V \setminus \mathcal{R}'')$ must be also conditionally Gibbs. Due to Theorem 6.5, the above algorithm is correct.

In [22], an algorithm is given for sampling hardcore model in a static setting, which can be expressed as our resampling meta-algorithm GenResample (Algorithm 3) with $\text{Expand}(\mathbf{X}, \mathcal{R}) = \text{vbl}(E^+(\mathcal{R}))$. Their algorithm can be interpreted as a static version of Algorithm 5. Because after the first iteration of Algorithm 5, the pair $(\mathbf{X}, \mathcal{R})$ must satisfy that $X_v = 1$ for all $v \in \mathcal{R}$, which means the $\text{Expand}(\mathbf{X}, \mathcal{R})$ in (24) is exactly $\text{vbl}(E^+(\mathcal{R}))$.

Proof of Theorem 4.6. For this algorithm, we define a new potential function $H_{\text{HC}} : 2^V \rightarrow \mathbb{Z}_{\geq 0}$ as

$$\forall \mathcal{R} \subseteq V : \quad H_{\text{HC}}(\mathcal{R}) \triangleq |E(\mathcal{R})|.$$

Let $\mathbf{X} \in \{0, 1\}^V$ be a configuration of hardcore model and $\mathcal{R} \subseteq V$ be a subset of vertices such that $(\mathbf{X}, \mathcal{R})$ is conditionally Gibbs with respect to the hardcore model \mathcal{I} . Fix one iteration of the while loop in Algorithm 5. The pair $(\mathbf{X}, \mathcal{R})$ is transformed to random pair $(\mathbf{X}', \mathcal{R}')$ after this iteration. Assume that $\lambda_v \leq \frac{1}{\sqrt{2\Delta}-1}$ for all $v \in V$. For any fixed \mathcal{R} and any fixed $X_{\mathcal{R}}$, we show the following holds

$$\mathbb{E}[H_{\text{HC}}(\mathcal{R}') \mid \mathcal{R}, X_{\mathcal{R}}] \leq \left(1 - \frac{1}{2\Delta}\right) H_{\text{HC}}(\mathcal{R}).$$

Since the above inequality holds for any fixed \mathcal{R} and $X_{\mathcal{R}}$, then we have

$$\mathbb{E}[H_{\text{HC}}(\mathcal{R}')] \leq \left(1 - \frac{1}{2\Delta}\right) \mathbb{E}[H_{\text{HC}}(\mathcal{R})].$$

Recall $\Delta = O(1)$. The value of the potential function decays with a constant factor in expectation. Therefore, Theorem 4.6 can be prove by going through the same proof as Theorem 4.1.

According to Algorithm 5, any edge $(u, v) \in E$ belongs to $E(\mathcal{R}')$ if and only if the values of X'_u and X'_v are both resampled as 1. Thus, we have

$$\forall (u, v) \in E : \quad (u, v) \in E(\mathcal{R}') \iff X'_u = X'_v = 1.$$

Therefore, to bound the expectation of $H_{\text{HC}}(\mathcal{R}') = |E(\mathcal{R}')|$, we can bound the probability of $X'_u = X'_v = 1$ for each edge $(u, v) \in E^+(\mathcal{R}'')$, where $\mathcal{R}'' = \text{Expand}(\mathbf{X}, \mathcal{R})$ is the expanded resample set. Recall \mathcal{R}'' depends only on \mathcal{R} and $X_{\mathcal{R}}$. Hence, the set \mathcal{R}'' is fixed.

Consider the edge $(u, v) \in E(\mathcal{R}'')$. Since X'_u and X'_v are sampled independently, then we have

$$\forall (u, v) \in E(\mathcal{R}'') : \quad \Pr[(u, v) \in E(\mathcal{R}')] = \Pr[X'_u = 1 \wedge X'_v = 1] = \frac{\lambda_u \lambda_v}{(1 + \lambda_u)(1 + \lambda_v)}. \quad (25)$$

Consider the edge $(u, v) \in \delta(\mathcal{R}'')$ where $u \in \mathcal{R}''$ and $v \notin \mathcal{R}''$. Since X'_u is sampled independently and $X'_v = X_v$, then we have

$$\forall (u, v) \in \delta(\mathcal{R}'') \text{ where } u \in \mathcal{R}'' \wedge v \notin \mathcal{R}'' : \quad \Pr[(u, v) \in E(\mathcal{R}')] = \frac{\lambda_u}{1 + \lambda_u} \Pr[X_v = 1]. \quad (26)$$

Note that the pair $(\mathbf{X}, \mathcal{R}'')$ is conditionally Gibbs with respect to \mathcal{I} . By the definition of Expand in (24), for any edge $(u, v) \in \delta(\mathcal{R}'')$ where $u \in \mathcal{R}''$ and $v \notin \mathcal{R}''$, it must hold that $X_u = 0$. Thus $X_{V \setminus \mathcal{R}''}$ is a random configuration sampled from the distribution of the hardcore model on induced subgraph $G[V \setminus \mathcal{R}'']$. We denote such distribution as $\mu_{V \setminus \mathcal{R}''}$. For any vertex $v \in V \setminus \mathcal{R}''$, suppose $\sigma \in \{0, 1\}^{V \setminus \mathcal{R}''}$ is an independent set on subgraph $G[V \setminus \mathcal{R}'']$ with $\sigma_v = 1$. Then $\sigma' \in \{0, 1\}^{V \setminus \mathcal{R}''}$ with $\sigma'_u = \sigma_u$ for $u \in V \setminus (\mathcal{R}'' \cup \{v\})$ and $\sigma'_v = 0$ is also an independent set on subgraph $G[V \setminus \mathcal{R}'']$. Note that $\frac{\mu_{V \setminus \mathcal{R}''}(\sigma)}{\mu_{V \setminus \mathcal{R}''}(\sigma')} = \lambda_v$. We have $\frac{\Pr[X_v=1]}{\Pr[X_v=0]} \leq \frac{\lambda_v}{1}$, which implies

$$\forall v \in V \setminus \mathcal{R}'' \quad \Pr[X_v = 1] \leq \frac{\lambda_v}{1 + \lambda_v}. \quad (27)$$

Let $\lambda = \max_{v \in V} \lambda_v$. Combining (25), (26) and (27), we have

$$\mathbb{E} [H_{\text{HC}}(\mathcal{R}') \mid \mathcal{R}, X_{\mathcal{R}}] \leq \sum_{e=(u,v) \in E^+(\mathcal{R}'')} \frac{\lambda_u \lambda_v}{(1 + \lambda_u)(1 + \lambda_v)} \leq \sum_{e=(u,v) \in E^+(\mathcal{R}'')} \left(\frac{\lambda}{1 + \lambda} \right)^2.$$

By the definition of Expand , it holds that $|E(\mathcal{R}'')| \leq (2\Delta - 1)|E(\mathcal{R})|$. This is because one edge at most incident to $2\Delta - 2$ edges in graph G , where Δ is the maximum degree of graph G . Similarly, the number of edges in $\delta(\mathcal{R}'')$ can be bounded as $|\delta(\mathcal{R}'')| \leq (2\Delta - 1)(\Delta - 1)|E(\mathcal{R})|$. Thus

$$\mathbb{E} [H_{\text{HC}}(\mathcal{R}') \mid \mathcal{R}, X_{\mathcal{R}}] \leq \Delta(2\Delta - 1) \left(\frac{\lambda}{1 + \lambda} \right)^2 H_{\text{HC}}(\mathcal{R}).$$

Since $\lambda_v \leq \frac{1}{\sqrt{2\Delta - 1}}$ for all $v \in V$, we have $\lambda \leq \frac{1}{\sqrt{2\Delta - 1}}$, which implies

$$\mathbb{E} [H_{\text{HC}}(\mathcal{R}') \mid \mathcal{R}, X_{\mathcal{R}}] \leq \Delta(2\Delta - 1) \left(\frac{1}{\sqrt{2\Delta}} \right)^2 H_{\text{HC}}(\mathcal{R}) \leq \left(1 - \frac{1}{2\Delta} \right) H_{\text{HC}}(\mathcal{R}).$$

The rest can be done by going through the same proof as Theorem 4.1. \square

9 A Lower Bound for Coloring

A proper q -coloring of a graph $G = (V, E)$ is a vector $\sigma \in [q]^V$ such that $\sigma_u \neq \sigma_v$ for all $(u, v) \in E$. The uniform distribution over all proper q -colorings of a graph $G = (V, E)$ can be expressed as the Gibbs distribution of a graphical model $\mathcal{I} = (V, E, [q], \Phi)$, where for each $v \in V$, ϕ_v is the uniform distribution over $[q]$, and for each $e \in E$, ϕ_e is the Boolean function indicating $\sigma_u \neq \sigma_v$.

Our dynamic sampler (Algorithm 1) is inefficient for sampling uniform proper q -coloring, in a sense that the only way to obtain a uniform proper q -coloring is to resample colors of all vertices independently and be very lucky to have this q -coloring happen to be proper. To see this, suppose that in one calling of Local-Resample (Algorithm 2) the current resample set is $\mathcal{R} \subseteq V$. Then all boundary edges $e \in \delta(\mathcal{R})$ must have $\kappa_e = 0$, which means that all boundary edges $e \in \delta(\mathcal{R})$ are going to join the new resample set \mathcal{R}' in the next iteration. Recall that the dynamic sampler stops when $\mathcal{R}' = \emptyset$, which can only happen when the resample set $\mathcal{R} = V$, but in such case the colors of all vertices are resampled independently and the chance of generating a proper q -coloring is exponentially small.

We further show that the above slow convergence might belong to a natural class of resampling algorithms. Consider any resampling procedure that transforms the random pair $(\mathbf{X}, \mathcal{R}) \in [q]^V \times 2^V$ of current coloring $\mathbf{X} \in [q]^V$ and resample set $\mathcal{R} \subseteq V$ to a new random pair $(\mathbf{X}', \mathcal{R}') \in [q]^V \times 2^V$. Assume that the followings hold for this resampling procedure:

- only the colors of vertices in \mathcal{R} are changed;
- the chain \mathfrak{M} on space $[q]^V \times 2^V$ with transition $(\mathbf{X}, V \setminus \mathcal{R}) \rightarrow (\mathbf{X}', V \setminus \mathcal{R}')$ as defined in Definition 4.2 satisfies the refined equilibrium condition as stated in Condition 2.

Then it holds that $\mathcal{R}' = \emptyset$ occurs only if $\mathcal{R} = V$. This means that we have to sample a proper q -coloring “from scratch”: the problem of sampling a uniform proper q -coloring is not made any easier by resampling.

To see why this is true, by contradiction, we assume that it is possible that $\mathcal{R} = R$ for some $R \neq V$ and $\mathcal{R}' = \emptyset$. There must exist an edge $(u, v) \in E$ with $v \in R$ and $u \notin R$. Let $(u, v) \in E$ be such an edge. Let $\sigma \in [q]^R$ be an assignment of colors to vertices in R . Suppose that we are conditioning on that $\mathcal{R} = R$ and $X_R = \sigma$. Since $(\mathbf{X}, V \setminus \mathcal{R})$ is conditionally Gibbs with respect to uniform proper q -coloring of graph G , we have $X_u \neq \sigma_v$. And because only variables in $\mathcal{R} = R$ are resampled, then we have $X'_u = X_u \neq \sigma_v$, which implies

$$\Pr[X'_u = \sigma_v] = 0.$$

Note that this guarantees that $\Pr[X'_u = \sigma_v] = 0$ conditioning on any particular \mathcal{R}' .

However, by the refined equilibrium condition, Condition 2, the new sample \mathbf{X}' is still conditionally Gibbs with respect to uniform proper q -coloring of graph G (even we start by conditioning on $\mathcal{R} = R$ and $X_R = \sigma$). In particular, conditioning on $\mathcal{R}' = \emptyset$, the new sample \mathbf{X}' is a uniformly distributed proper q -coloring, which means

$$\Pr[X'_u = \sigma_v] = \frac{1}{q}.$$

A contradiction.

10 Conclusion and Future Work

We give a dynamic sampling method that allows us to sample perfectly from a broad class of graphical models, while variables and constraints of the graphical model are changing dynamically. We provide sufficient conditions under which such algorithms run incrementally in time proportional to the size of the update. A key to these results is to establish certain equilibrium condition satisfied by local resampling. On specific graphical models, this equilibrium condition also helps to obtain better convergence of the algorithm.

For graphical models with “truly repulsive” hard constraints, e.g. uniform proper q -coloring, our current technique meets certain barrier. Very recently, a new technique of *dynamic MCMC sampling* [8] was invented which uses a data structure of super-linear size to dynamically maintain a sample. In contrast, the technique of this paper maintains a sample with both small space cost and small incremental time cost.

One open problem is on specific graphical models, to improve the regimes for efficient dynamic sampling to the uniqueness regimes. For example, for the hardcore model, such result would imply a highly efficient (near-linear time for graphs with bounded maximum degree and near-quadratic time for general graphs) algorithm for sampling from the hardcore model in the uniqueness regime, which remains to be open even for static and approximate sampling. So far such efficient static and approximate sampling algorithms are only known for graphs with large girth and sufficiently large degree [7]. Poly-time algorithms (with very large polynomial) were known for graphs with bounded maximum degree [50] and only very recently for general graphs [1].

Along this direction, a very interesting open problem is to give dynamic samplers from Gibbs distributions with mild decay of correlation. One major open problem is dynamically sampling uniform matchings, which has a decay of correlation with rate $1 - O(1/\sqrt{\Delta})$ [2]. A fast dynamic sampler for matchings with, say $O(\Delta^{1.5})$ incremental cost per each update of an edge, even being used as a static and approximate sampler, would improve the $\tilde{O}(n^2m)$ time bound of the Jerrum-Sinclair chain for matchings [32].

More broadly, our techniques should be of independent interest and, in particular, should be useful to extend our results to sampling from joint distributions over continuous distributions and/or with global constraints.

Acknowledgement

This research is supported by the National Key R&D Program of China 2018YFB1003202 and the National Science Foundation of China under Grant Nos. 61722207 and 61672275. The authors would like to acknowledge the hospitality of the Bernoulli Center at EPFL where part of this work was done.

References

- [1] N. ANARI, K. LIU, AND S. O. GHARAN, *Spectral independence in high-dimensional expanders and applications to the hardcore model*, arXiv preprint arXiv:2001.00303, (2020).

- [2] M. BAYATI, D. GAMARNIK, D. KATZ, C. NAIR, AND P. TETALI, *Simple deterministic approximation algorithms for counting matchings*, in Proceedings of the 39th Annual ACM Symposium on Theory of Computing (STOC), 2007, pp. 122–127.
- [3] C. M. BISHOP, *Pattern Recognition and Machine Learning*, Springer, 2006.
- [4] A. DEMBO AND A. MONTANARI, *Ising models on locally tree-like graphs*, The Annals of Applied Probability, 20 (2010), pp. 565–592.
- [5] C. DEMETRESCU, D. EPPSTEIN, Z. GALIL, AND G. F. ITALIANO, *Dynamic graph algorithms*, in Algorithms and Theory of Computation Handbook, Second Edition, Volume 1, Chapman and Hall/CRC, 2009, pp. 228–255.
- [6] M. DYER AND C. GREENHILL, *On markov chains for independent sets*, Journal of Algorithms, 35 (2000), pp. 17–49.
- [7] C. EFTHYMIU, T. P. HAYES, D. ŠTEFANKOVIĆ, E. VIGODA, AND Y. YIN, *Convergence of MCMC and loopy BP in the tree uniqueness region for the hard-core model*, in Proceedings of the 57th Annual IEEE Symposium on Foundations of Computer Science (FOCS), 2016, pp. 704–713.
- [8] W. FENG, K. HE, X. SUN, AND Y. YIN, *Dynamic mcmc sampling*, arXiv preprint arXiv:1904.11807, (2019).
- [9] W. FENG, Y. SUN, AND Y. YIN, *What can be sampled locally?*, in Proceedings of the 36th ACM Symposium on Principles of Distributed Computing (PODC), 2017, pp. 121–130.
- [10] J. A. FILL, *An interruptible algorithm for perfect sampling via Markov chains*, in Proceedings of the 29th Annual ACM Symposium on Theory of Computing (STOC), 1997, pp. 688–695.
- [11] J. A. FILL AND M. HUBER, *The randomness recycler: a new technique for perfect sampling*, in Proceedings of the 41st Annual IEEE Symposium on Foundations of Computer Science (FOCS), 2000, pp. 503–511.
- [12] J. A. FILL, M. MACHIDA, D. J. MURDOCH, AND J. S. ROSENTHAL, *Extension of Fill’s perfect rejection sampling algorithm to general chains*, Random Structures & Algorithms, 17 (2000), pp. 290–316.
- [13] M. FISCHER AND M. GHAFARI, *A simple parallel and distributed sampling technique: Local glauber dynamics*, in International Symposium on Distributed Computing (DISC), 2018.
- [14] A. GALANIS, D. ŠTEFANKOVIČ, AND E. VIGODA, *Inapproximability for antiferromagnetic spin systems in the tree nonuniqueness region*, Journal of the ACM (JACM), 62 (2015), p. 50.
- [15] A. GALANIS, D. ŠTEFANKOVIČ, AND E. VIGODA, *Inapproximability of the partition function for the antiferromagnetic Ising and hard-core models*, Combinatorics, Probability and Computing, 25 (2016), pp. 500–559.
- [16] A. GERSCHCNFELD AND A. MONIANARI, *Reconstruction for models on random graphs*, in Proceedings of the 48th Annual IEEE Symposium on Foundations of Computer Science (FOCS), 2007.

- [17] L. GOLDBERG, R. MARTIN, AND M. PATERSON, *Strong spatial mixing with fewer colors for lattice graphs*, SIAM Journal on Computing, 35 (2005), p. 486.
- [18] H. GUO AND K. HE, *Tight bounds for popping algorithms*, arXiv preprint arXiv:1807.01680, (2018).
- [19] H. GUO AND M. JERRUM, *Perfect simulation of the hard disks model by partial rejection sampling*, in Proceedings of the 45th International Colloquium on Automata, Languages, and Programming (ICALP), 2018.
- [20] H. GUO AND M. JERRUM, *A polynomial-time approximation algorithm for all-terminal network reliability*, in Proceedings of the 45th International Colloquium on Automata, Languages, and Programming (ICALP), 2018.
- [21] H. GUO AND M. JERRUM, *Random cluster dynamics for the Ising model is rapidly mixing*, The Annals of Applied Probability, 28 (2018), pp. 1292–1313.
- [22] H. GUO, M. JERRUM, AND J. LIU, *Uniform sampling through the Lovász local lemma*, in Proceedings of the 49th Annual ACM Symposium on Theory of Computing (STOC), 2017, pp. 342–355.
- [23] B. HAEUPLER, B. SAHA, AND A. SRINIVASAN, *New constructive aspects of the Lovász local lemma*, Journal of the ACM (JACM), 58 (2011), p. 28.
- [24] D. G. HARRIS, *New bounds for the Moser-Tardos distribution*, arXiv preprint arXiv:1610.09653, (2016).
- [25] D. G. HARRIS AND A. SRINIVASAN, *The Moser-Tardos framework with partial resampling*, in Proceedings of the 54th Annual IEEE Symposium on Foundations of Computer Science (FOCS), 2013, pp. 469–478.
- [26] N. J. HARVEY, P. SRIVASTAVA, AND J. VONDRÁK, *Computing the independence polynomial: from the tree threshold down to the roots*, in Proceedings of the 29th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), SIAM, 2018, pp. 1557–1576.
- [27] N. J. HARVEY AND J. VONDRÁK, *An algorithmic proof of the Lovász local lemma via re-sampling oracles*, in Proceedings of the 56th Annual IEEE Symposium on Foundations of Computer Science (FOCS), 2015, pp. 1327–1346.
- [28] T. P. HAYES AND A. SINCLAIR, *A general lower bound for mixing of single-site dynamics on graphs*, The Annals of Applied Probability, (2007), pp. 931–952.
- [29] M. HUBER, *Perfect sampling using bounding chains*, The Annals of Applied Probability, 14 (2004), pp. 734–753.
- [30] M. L. HUBER, *Perfect simulation*, Chapman and Hall/CRC, 2016.
- [31] M. JERRUM, *A very simple algorithm for estimating the number of k -colorings of a low-degree graph*, Random Structures & Algorithms, 7 (1995), pp. 157–165.

- [32] M. JERRUM, *Counting, sampling and integrating: algorithms and complexity*, Springer Science & Business Media, 2003.
- [33] M. JERRUM AND A. SINCLAIR, *Polynomial-time approximation algorithms for the Ising model*, SIAM Journal on Computing, 22 (1993), pp. 1087–1116.
- [34] M. JERRUM, A. SINCLAIR, AND E. VIGODA, *A polynomial-time approximation algorithm for the permanent of a matrix with nonnegative entries*, Journal of the ACM (JACM), 51 (2004), pp. 671–697.
- [35] M. JERRUM, L. G. VALIANT, AND V. V. VAZIRANI, *Random generation of combinatorial structures from a uniform distribution*, Theoretical Computer Science, 43 (1986), pp. 169–188.
- [36] M. I. JORDAN, *Learning in graphical models*, vol. 89, Springer Science & Business Media, 1998.
- [37] D. E. KNUTH, *The Art of Computer Programming, Volume 2 (3rd Ed.): Seminumerical Algorithms*, Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 1997.
- [38] D. KOLLER, N. FRIEDMAN, AND F. BACH, *Probabilistic graphical models: principles and techniques*, MIT press, 2009.
- [39] D. A. LEVIN AND Y. PERES, *Markov chains and mixing times*, vol. 107, American Mathematical Soc., 2017.
- [40] M. MEZARD AND A. MONTANARI, *Information, physics, and computation*, Oxford University Press, 2009.
- [41] A. MOITRA, *Approximate counting, the lovasz local lemma, and inference in graphical models*, in Proceedings of the 49th Annual ACM Symposium on Theory of Computing (STOC), ACM, 2017, pp. 356–369.
- [42] R. A. MOSER AND G. TARDOS, *A constructive proof of the general Lovász local lemma*, Journal of the ACM (JACM), 57 (2010), p. 11.
- [43] E. MOSSEL AND A. SLY, *Exact thresholds for Ising–Gibbs samplers on general graphs*, The Annals of Probability, 41 (2013), pp. 294–328.
- [44] J. G. PROPP AND D. B. WILSON, *Exact sampling with coupled Markov chains and applications to statistical mechanics*, Random Structures & Algorithms, 9 (1996), pp. 223–252.
- [45] J. SALAS AND A. D. SOKAL, *Absence of phase transition for antiferromagnetic Potts models via the Dobrushin uniqueness theorem*, Journal of Statistical Physics, 86 (1997), pp. 551–579.
- [46] D. STRASZAK AND N. K. VISHNOI, *Real stable polynomials and matroids: Optimization and counting*, in Proceedings of the 49th Annual ACM Symposium on Theory of Computing (STOC), 2017, pp. 370–383.
- [47] E. VIGODA, *Fast convergence of the glauber dynamics for sampling independent sets: Part II*, Tech. Report TR-99-003, International Computer Science Institute, 1999.

- [48] E. VIGODA, *Improved bounds for sampling colorings*, Journal of Mathematical Physics, 41 (2000), pp. 1555–1569.
- [49] M. J. WAINWRIGHT AND M. I. JORDAN, *Graphical models, exponential families, and variational inference*, Foundations and Trends in Machine Learning, 1 (2008), pp. 1–305.
- [50] D. WEITZ, *Counting independent sets up to the tree threshold*, in Proceedings of the 38th Annual ACM Symposium on Theory of Computing (STOC), 2006, pp. 140–149.
- [51] D. B. WILSON, *Generating random spanning trees more quickly than the cover time*, in Proceedings of the 28th Annual ACM Symposium on Theory of Computing (STOC), 1996, pp. 296–303.
- [52] J. S. YEDIDIA, W. T. FREEMAN, AND Y. WEISS, *Constructing free-energy approximations and generalized belief propagation algorithms*, IEEE Transactions on Information Theory, 51 (2005), pp. 2282–2312.