

RESEARCH ARTICLE

WILEY

Fast algorithms at low temperatures via Markov chains[†]

Zongchen Chen¹ | Andreas Galanis² | Leslie A. Goldberg² | Will Perkins³ | James Stewart² | Eric Vigoda¹

¹School of Computer Science, Georgia Institute of Technology, Atlanta, Georgia

²Department of Computer Science, University of Oxford, Oxford, United Kingdom

³Department of Mathematics, Statistics, and Computer Science, University of Illinois at Chicago, Chicago, Illinois

Correspondence

James Stewart, School of Computer Science, Georgia Institute of Technology, Atlanta, GA.
Email: james.stewart@chch.ox.ac.uk

Funding information

This research was supported by the NSF grants, CCF-1617306, CCF-1563838 (Z.C.), DMS-1847451, CCF-1934915 (W.P.), European Research Council under the European Union's Seventh Framework Programme (FP7/2007-2013) ERC grant, Agreement no. 334828 (A.G.).

Abstract

Efficient algorithms for approximate counting and sampling in spin systems typically apply in the so-called high-temperature regime, where the interaction between neighboring spins is “weak.” Instead, recent work of Jenssen, Keevash, and Perkins yields polynomial-time algorithms in the low-temperature regime on bounded-degree (bipartite) expander graphs using polymer models and the cluster expansion. In order to speed up these algorithms (so the exponent in the run time does not depend on the degree bound) we present a Markov chain for polymer models and show that it is rapidly mixing under exponential decay of polymer weights. This yields, for example, an $O(n \log n)$ -time sampling algorithm for the low-temperature ferromagnetic Potts model on bounded-degree expander graphs. Combining our results for the hard-core and Potts models with Markov chain comparison tools, we obtain polynomial mixing time for Glauber dynamics restricted to appropriate portions of the state space.

KEYWORDS

approximate counting, expander graphs, hard-core model, Markov chains, Potts model

[†]These results were announced in preliminary form (without proofs) as a brief abstract in the proceedings of APPROX/RANDOM 2019.

This is an open access article under the terms of the Creative Commons Attribution License, which permits use, distribution and reproduction in any medium, provided the original work is properly cited.

© 2020 The Authors. *Random Structures & Algorithms* published by Wiley Periodicals LLC.

1 | INTRODUCTION

The hard-core model from statistical physics is defined on the set of independent sets of a graph G , where the independent sets are weighted by a fugacity $\lambda > 0$. The associated Gibbs distribution $\mu_{G,\lambda}$ is defined as follows, for an independent set I :

$$\mu_{G,\lambda}(I) = \frac{\lambda^{|I|}}{Z_{G,\lambda}}, \quad (1)$$

where $Z_{G,\lambda} = \sum_{I \in \mathcal{I}(G)} \lambda^{|I|}$ is the hard-core partition function (also called the independence polynomial), $\mathcal{I}(G)$ is the set of independent sets of G , and $\lambda > 0$ is the *fugacity*.

In applications, there are two important computational tasks associated to a spin model such as the hard-core model. Given an error parameter $\varepsilon \in (0, 1)$, an ε -approximate counting algorithm outputs a number \hat{Z} so that $e^{-\varepsilon} Z_{G,\lambda} \leq \hat{Z} \leq e^{\varepsilon} Z_{G,\lambda}$, and an ε -approximate sampling algorithm outputs a random sample I with distribution $\hat{\mu}$ so that the total variation distance satisfies $\|\mu_{G,\lambda} - \hat{\mu}\|_{TV} < \varepsilon$.

While classical statistical physics is most interested in studying the hard-core model on the integer lattice \mathbb{Z}^d , the perspective of computer science is to consider wider families of graphs, such as the set of all graphs, all graphs of maximum degree Δ , or all bipartite graphs of maximum degree Δ .

Almost all proven efficient algorithms for approximate counting and sampling from the hard-core model work for low fugacities (the weak interaction regime, akin to the low temperature regime of the Potts model). In the high temperature regime there are at least three distinct algorithmic approaches to approximate counting and sampling: Markov chains, correlation decay, and polynomial interpolation. One striking advantage of the Markov chain approach is that the algorithms are much faster and simpler than the algorithms from the other approaches. In particular, it is common for a Markov chain sampling algorithm to run in time $O(n \log n)$, for example, see [13, 14], while typical running times for algorithms based on correlation decay [26, 32] and polynomial interpolation [1] are $n^{O(\log \Delta)}$ where Δ is the maximum degree of the graph.

In general there are no known efficient algorithms at low temperatures (high fugacities), but recently efficient algorithms have been developed for some special classes of graphs including subsets of \mathbb{Z}^d [18], random regular bipartite graphs, and bipartite expander graphs in general [21, 25]. What these bipartite graphs have in common is that for large enough λ , typical independent sets drawn from $\mu_{G,\lambda}$ align closely with one side or the other of the bipartition (the two ground states). This phenomenon is related to the phase transition phenomenon in infinite graphs, and implies the exponentially slow mixing time of local Markov chains [5, 16, 27]. The algorithms introduced in [18] exploit this phenomenon by expressing the partition function $Z_{G,\lambda}$ in terms of deviations from the two ground states, and then using a truncation of a convergent series expansion (the Taylor series or the cluster expansion) to approximate the log partition function. In statistical physics this is called a *perturbative* approach, and while in general it does not work in the largest possible range of parameter space, when it does work it gives a very detailed probabilistic understanding of the model [7, 10, 29].

To apply the so-called perturbative approach at low temperatures, one rewrites the original spin model as a new model in which single spin interactions are replaced by the interaction of connected components representing deviations from a chosen ground state. Such models are known in general as *abstract polymer models* [23], see Section 1.1 for the polymer models we consider here, and have long been used in statistical physics to understand phase transitions. In this paper, we show that once a low-temperature spin model has been transformed into a polymer model, Markov chains once again become an effective algorithmic tool. Using this approach we obtain nearly linear and quadratic time

sampling algorithms for low temperature models on expander graphs in cases where only $n^{O(\log \Delta)}$ -time algorithms were previously known.

1.1 | Subset polymer models

Abstract polymer models, as defined by Kotecký and Preiss [23], are an important tool in studying the equilibrium phases of statistical physics models on lattices, see, for example, [7, 24] among many others.¹ The paper [4] has a more detailed history of their use in statistical physics and combinatorics. Recently, polymer models have been used to develop efficient deterministic algorithms for sampling and approximating the partition functions of statistical physics models on lattices [18] and expander graphs [21, 25] at low temperatures, the regime in which Markov chains like the Glauber dynamics are known to mix slowly.

We will study the following class of abstract polymer models, known as subset polymer models (defined by Gruber and Kunz [17]). We begin by describing the relevant polymers: start with a finite host graph G and a set $[q] = \{0, \dots, q-1\}$ of spins. For each vertex v , there is a ground-state spin g_v . A polymer γ consists of a connected set of vertices together with an assignment σ_γ of spins from $\{0, \dots, q-1\} \setminus g_v$ to each vertex $v \in \gamma$ (we abuse notation and use γ to denote both the polymer and the associated set of vertices). The size of a polymer, $|\gamma|$, is the number of vertices in γ . The set of all polymers is $\mathcal{P}(G)$.

A polymer model on G consists of a set $\mathcal{C}(G) \subseteq \mathcal{P}(G)$ of “allowed” polymers, and a nonnegative weight w_γ for each polymer $\gamma \in \mathcal{C}(G)$. We denote this model by $(\mathcal{C}(G), w)$. Two polymers γ and γ' are called “compatible” (written $\gamma \sim \gamma'$) if their distance in the host graph is at least 2; otherwise they are “incompatible” (written $\gamma \not\sim \gamma'$). The state space of allowable configurations is $\Omega = \{\Gamma \subseteq \mathcal{C}(G) \mid \forall \gamma, \gamma' \in \Gamma, \gamma \sim \gamma'\}$.

The partition function of the polymer model is

$$Z(G) = \sum_{\Gamma \in \Omega} \prod_{\gamma \in \Gamma} w_\gamma,$$

where the empty set of polymers contributes 1 to the partition function. The Gibbs measure μ_G is the probability distribution on Ω given by

$$\mu_G(\Gamma) = \frac{\prod_{\gamma \in \Gamma} w_\gamma}{Z(G)}.$$

Note that the polymer model is in fact a hard-core model on the “incompatibility graph” of $\mathcal{C}(G)$, where two polymers are joined by an edge if they are incompatible, with nonuniform fugacities given by the weights w_γ . The geometry inherited from the host graph G and the sizes of the polymers adds additional structure to the model.

Example 1. One instance of a polymer model is the hard-core model itself: polymers are single vertices of the graph G , labeled with “1” (for occupied) against a ground state “0” (for unoccupied). Each polymer (vertex) v comes with the weight function $w_v = \lambda$. Then the set of allowable polymer configurations is exactly the set of independent sets of G , and so the polymer model partition function is exactly the partition function of the hard-core model on G .

¹See also the relevant notion of “animal models” by Dobrushin [10].

Example 2. A second instance of a polymer model is related to the ferromagnetic q -color Potts model on a graph G (see Definition 8). Fix a color $g \in [q]$ to be the ground state color, and define polymers to be connected subgraphs of G of size at most M , with vertices labeled by the remaining colors $[q] \setminus \{g\}$. A polymer γ has weight function $w_\gamma = e^{-\beta B(\gamma)}$ where $B(\gamma)$ is the number of bichromatic edges in γ plus the size of the edge boundary of γ in G . A configuration of compatible polymers maps to a unique Potts configuration σ in which all connected components of non- g -colored vertices have size at most M , and the weight of σ in the Potts model is exactly the product of the weight functions of the polymers. The polymer model partition function $Z(G)$, with an appropriate choice of M , represents the contribution to the Potts model partition function of colorings where color g “dominates,” see also Section 1.3 for more details.

As with the hard-core model, there are two main computational problems associated to a polymer model: approximate sampling from μ_G and approximate counting of $Z(G)$. We will approach them both via Markov chain algorithms. In general we will be interested in families of polymer models defined on classes of graphs. We denote such a family $(C(\cdot), w, \mathcal{G})$, where for each graph $G \in \mathcal{G}$, $(C(G), w)$ is a polymer model. We will always use n to denote the number of vertices of a graph G .

We consider two conditions on the weight functions w_γ and give their algorithmic consequences.

Definition 1. A polymer model $(C(\cdot), w, \mathcal{G})$ satisfies the *polymer mixing condition* if there exists $\theta \in (0, 1)$ such that

$$\sum_{\gamma' \sim \gamma} |\gamma'| w_{\gamma'} \leq \theta |\gamma| w_\gamma \quad (2)$$

for all $G \in \mathcal{G}$ and all $\gamma \in C(G)$.

We postpone the formal definition of mixing time to Section 2.2 and state our first main result here.

Theorem 2. Suppose that a polymer model $(C(\cdot), w, \mathcal{G})$ satisfies the polymer mixing condition (2). Then for each $G \in \mathcal{G}$ there is a Markov chain making single polymer updates with stationary distribution μ_G and mixing time $T_{\text{mix}}(\epsilon) = O(n \log(n/\epsilon))$.

Theorem 2 on its own does not guarantee an efficient algorithm for sampling from μ_G because the Markov chain only yields an efficient sampling algorithm if we can implement each step efficiently. We will show that under a stronger condition we can do this.

Definition 3. A polymer model $(C(\cdot), w, \mathcal{G})$ is said to be *computationally feasible* if, for each $G \in \mathcal{G}$ and each $\gamma \in \mathcal{P}(G)$, we can determine, in time polynomial in $|\gamma|$, whether $\gamma \in C(G)$, and compute w_γ if it is.

Definition 4. A polymer model $(C(\cdot), w, \mathcal{G})$ with q spins on a class \mathcal{G} of graphs of maximum degree Δ satisfies the *polymer sampling condition* with constant $\tau \geq 5 + 3 \log((q-1)\Delta)$ if

$$w_\gamma \leq e^{-\tau |\gamma|}. \quad (3)$$

for all $G \in \mathcal{G}$ and all $\gamma \in C(G)$.

We have the following theorem.

Theorem 5. *If a computationally feasible polymer model $(C(\cdot), w, \mathcal{G})$ satisfies the polymer sampling condition (3) then for all $G \in \mathcal{G}$ there is an ε -approximate sampling algorithm for μ_G with running time $O(n \log(n/\varepsilon))$.*

Note that the polymer sampling condition required by Theorem 5 is in general more demanding than the zero-freeness required by cluster-expansion algorithms, but, as Theorem 5 demonstrates, it leads to faster algorithms. Finally, we can use the sampling algorithm and simulated annealing to give a fully polynomial time randomized approximation scheme (FPRAS) for computing the partition function of polymer models.

Theorem 6. *If a computationally feasible polymer model $(C(\cdot), w, \mathcal{G})$ satisfies the polymer sampling condition (3) then for all $G \in \mathcal{G}$ there is a randomized ε -approximate counting algorithm for $Z(G)$ with running time $O((n/\varepsilon)^2 \log^2(n/\varepsilon))$ and success probability at least $3/4$.*

Fernández, Ferrari, and Garcia [15] introduced a condition very similar to the polymer mixing condition in the setting of polymer models on \mathbb{Z}^d . Their objective was to derive probabilistic properties of polymer models directly, without going through the combinatorics and complex analysis inherent in the cluster expansion for the log partition function. They introduced a continuous time stochastic process whose stationary distribution was the infinite volume Gibbs measure of their polymer model and their version of condition (2) implied an exponentially fast rate of convergence of this process. They remarked that such an approach had the potential to be an efficient computational tool.

Here we take an algorithmic point of view, and use the polymer mixing and sampling conditions to show that a simple discrete time Markov chain mixes rapidly and can be used to design efficient sampling and approximation algorithms. Our approach differs from that of [15] in that while they are interested primarily in the probabilistic properties of spin models on \mathbb{Z}^d , we are interested in algorithmic problems involving spin models on general families of graphs. Our setting of discrete time processes on finite graphs is also more suitable to studying algorithmic questions. Our work confirms the central point of [15]: that complex analysis and absolute convergence of the cluster expansion is not necessary to derive many important properties of a polymer model.

1.2 | Applications

We apply our results for subset polymer models to two specific examples: the ferromagnetic Potts model and the hard-core model on expander graphs. To state these results we need some definitions.

Definition 7. Let $\alpha > 0$. A graph G is an α -expander graph if for all $S \subset V(G)$ with $|S| \leq |V(G)|/2$, we have $e(S, S^c) \geq \alpha|S|$, where $S^c = V(G) \setminus S$ and $e(S, S^c)$ is the number of edges exiting the set S .

Definition 8. The q -color ferromagnetic Potts model with parameter $\beta > 0$ is a random assignment of q colors to the vertices of a graph defined by

$$\mu_{G,\beta}(\sigma) = \frac{e^{-\beta m(G,\sigma)}}{Z_{G,\beta}},$$

where $m(G, \sigma)$ is the number of bichromatic edges of G under the coloring σ and $Z_{G,\beta} = \sum_{\sigma} e^{-\beta m(G,\sigma)}$ is the Potts model partition function. The parameter β is known as the inverse temperature.

Jenssen, Keevash, and Perkins [21] gave an FPTAS and polynomial-time sampling algorithm for the Potts model on expander graphs, with an algorithm based on the cluster expansion and Barvinok's method of polynomial interpolation. Under essentially the same conditions on the parameters we give a Markov chain based sampling algorithm with near linear running time.

Theorem 9. *Suppose $q \geq 2$, $\Delta \geq 3$ are integers and $\alpha > 0$ is a real. Then for $\beta \geq \frac{5+3 \log((q-1)\Delta)}{\alpha}$ and any $qe^{-n} \leq \epsilon < 1$, there is an ϵ -approximate sampling algorithm for the q -state ferromagnetic Potts model with parameter β on all n -vertex α -expander graphs of maximum degree Δ with running time $O(n \log(n/\epsilon))$. There is also an ϵ -approximate counting algorithm with running time $O((n/\epsilon)^2 \log^2(n/\epsilon))$ and success probability at least $3/4$.*

Note that, if the desired error satisfies $\epsilon < qe^{-n}$, then we can simply compute the partition function by brute force in $\text{poly}(n/\epsilon)$ time. This observation combined with the above result gives an FPRAS, but we can no longer guarantee a running time of $O((n/\epsilon)^2 \log^2(n/\epsilon))$ for exponentially small values of ϵ . A similar point also applies to the algorithm that we give for the hard-core model.

Definition 10. Let $\alpha \in (0, 1)$. A bipartite graph $G = (V, E)$ with bipartition $V = V^0 \cup V^1$ is a bipartite α -expander if, for $i \in \{0, 1\}$ and all $S \subseteq V^i$ where $|S| \leq |V^i|/2$, we have $N_G(S) \geq (1 + \alpha)|S|$ where $N_G(S)$ denotes the set of vertices that are adjacent to some vertex in S .

Again we give a fast Markov chain based algorithm for sampling from the hard-core model for essentially the same range of parameters for which an FPTAS is given in [21].

Theorem 11. *Suppose $\Delta \geq 3$ is an integer and $\alpha \in (0, 1)$ is a real. Then for any $\lambda \geq (3\Delta)^{6/\alpha}$ and $4e^{-n} \leq \epsilon < 1$, there is an ϵ -approximate sampling algorithm for the hard-core model with parameter λ on all n -vertex bipartite α -expander graphs of maximum degree Δ . There is also an ϵ -approximate counting algorithm for the hard-core model with success probability at least $1 - \epsilon$. Both algorithms run in time $O((n/\epsilon)^2 \log^3(n/\epsilon))$.*

The extra factor of n in the running time of the sampling algorithm for the hard-core model as compared to the Potts model is due to the fact that the hard-core model on a bipartite graph does not in general exhibit exact symmetry between the ground states, and so we must approximate the partition functions of the even and odd dominant independent sets to sample.

We can extend these algorithms to obtain fast sampling algorithms in most situations in which a counting problem can be put in the framework of subset polymer models. For instance, we can use Theorems 5 and 6 to improve the running times of the algorithms given by [22, 25] for sampling and counting proper q -colorings in Δ -regular bipartite graphs (for large Δ). The two papers give slightly different polymer models for proper q -colorings on Δ -regular bipartite graphs—see [22, Section 5] and [25, Section 5.2]. Section 5.2 of [25] shows that their polymer model is computationally feasible. Section 5.1 of [22] shows that their polymer model satisfies the Kotecký-Preiss condition—in fact, their proof establishes the polymer sampling condition (3). It is easy to see (by comparing the polymer weights) that the polymer model of [25] therefore also satisfies the polymer sampling condition. Thus, we get the following corollary of Theorem 5 and 6.

Corollary 12. *There is an absolute constant $C > 0$ so that for all even $q \geq 3$, all $\Delta \geq Cq^2 \log^2 q$ and all $\epsilon > e^{-n/(8q)}$, there is an ϵ -approximate sampling algorithm to sample a uniformly random proper q -coloring from a random Δ -regular bipartite graph running in time $O(n \log(n/\epsilon))$. Furthermore, there*

is a randomized ε -approximation algorithm for the number of proper q -colorings with running time $O((n/\varepsilon)^2 \log^2(n/\varepsilon))$ and success probability at least $3/4$. For odd q , there are ε -approximate counting and sampling algorithms that both run in time $O((n/\varepsilon)^2 \log^3(n/\varepsilon))$.

As with independent sets, the extra factor of n in the running time for odd q comes from the fact that the ground states (colorings in which one side of the bipartition is assigned $\lfloor q/2 \rfloor$ colors and the other side $\lfloor q/2 \rfloor$ colors) are exactly symmetric only if q is even.

Finally, we remark that the approximate counting algorithms for these applications based on truncating the cluster expansion can run faster than $n^{O(\log \Delta)}$ if the parameters (expansion, fugacity, inverse temperature) are high enough (see [22], Theorem 8), but the sampling algorithms derived from this approach will not match the $\tilde{O}(n)$ or $\tilde{O}(n^2)$ sampling algorithms we obtain here.

1.3 | Comparison to spin Glauber dynamics

A very natural idea to sample at low temperatures (large β for the Potts model, large λ for the hard-core model) is to use a single-spin update Markov chain like the Glauber dynamics, but to start in one of the ground states of the model chosen at random. For example, pick one of the q -colors with equal probability then start the Potts model Glauber dynamics in the monochromatic configuration with that color. The intuition is that the Glauber dynamics will mix well within the portion of the state space close to the chosen ground state, and the randomness in the choice of ground state will ensure that an accurate sample from the full measure is obtained. Analyzing this algorithm was suggested in [18, 21].

While we are not yet able to show that this algorithm succeeds, we make partial progress. We show that Glauber dynamics, restricted to remain in a portion of the state space, mixes rapidly (in polynomial time). It is easiest to state our result for the ferromagnetic Potts model.

For a ground state color $g \in [q]$ and an integer M , let $\Omega_M^g(G)$ be the set of q -colorings of the vertices of G so that every connected component of G colored with the palette of colors $[q] \setminus g$ is of size at most M . The set $\Omega_M^g(G)$ consists of colorings that come from the valid polymer configurations from Example 2. In [21] it is shown that for an appropriate choice of M , the set $\{\Omega_M^g(G)\}_{g \in [q]}$ forms an “almost partition” of the set of all colorings, in that the weight of both the overlap of the almost partition and the set of colorings uncovered by the almost partition is at most ε under the conditions of Theorem 9. In particular, an ε -approximate sample from the Potts model restricted to $\Omega_M^g(G)$ for $M = O(\log(n/\varepsilon))$ is enough (by symmetry) to obtain a $(q\varepsilon)$ -approximate sample from the Potts distribution $\mu_{G,\beta}$ (cf. Lemma 28). Using Markov chain comparison, we show in Section 5.3.1 that this can be done using the usual spin Glauber dynamics restricted to remain in $\Omega_M^g(G)$.

Theorem 13. *Suppose $q \geq 2$, $\Delta \geq 3$ are integers and $\alpha > 0$ is a real. Let $\beta \geq \frac{5+3 \log((q-1)\Delta)}{\alpha}$ be a real number and $g \in [q]$. Then, for any n -vertex α -expander graph G of maximum degree Δ and any $\varepsilon \in (0, 1)$, for $M = O(\log(n/\varepsilon))$ the Glauber dynamics restricted to $\Omega_M^g(G)$ has mixing time $T_{\text{mix}}(\varepsilon)$ polynomial in n and $1/\varepsilon$.*

We remark that the polynomial bound in Theorem 13 depends on q, Δ, α exponentially, see the relevant Theorem 25 and Section 5.3.1 for details. Theorem 13 shows that despite exponentially slow mixing of the Glauber dynamics on the full state space [3], it can still be used to obtain a polynomial-time approximate sampling algorithm. We leave for future work two important extensions that would complete the picture: (1) showing that *unrestricted* Glauber dynamics starting from a well-chosen configuration works and (2) lowering the running time to $O(n \log n)$ from the large polynomial we obtain in the theorem.

In Section 5, we state a general theorem (Theorem 25) comparing the polymer model dynamics to spin model dynamics as well as a specific result for the hard-core model (Theorem 29).

2 | POLYMER MODELS AND MARKOV CHAINS

Here we compare various conditions on the weight functions of a polymer model, namely the Kotecký-Preiss [23] condition and the polymer sampling condition, and show that the latter implies the former. Then, we define the polymer Markov chain which we use to prove Theorems 2 and 5.

2.1 | A comparison of the conditions on the weights

Here we show that the polymer sampling condition (3) implies the well-known Kotecký-Preiss [23] condition:

$$\sum_{\gamma' \sim \gamma} e^{|\gamma'|} w_{\gamma'} \leq |\gamma|.$$

To see the implication, we use a lemma of Borgs, Chayes, Kahn, and Lovász.

Lemma 14 ([6]). *Let G have maximum degree $\Delta \geq 3$ and let $v \in V(G)$. The number of connected induced subgraphs of G of size k containing v is at most $(e\Delta)^{k-1}$.*

Now consider a polymer model satisfying (3) with constant $\tau \geq 5 + 3 \log((q-1)\Delta)$. Fix $\gamma \in C(G)$. We have that

$$\sum_{\gamma' \sim \gamma} e^{|\gamma'|} w_{\gamma'} \leq \sum_{v \in \gamma \cup \partial \gamma} \sum_{k \geq 1} \sum_{\gamma' \in C(G); |\gamma'|=k, v \in \gamma'} e^k e^{-\tau k}.$$

In order to account for all of the polymers that we sum over in the above, we consider the connected induced subgraphs of G of size k that contain v , and the assignments to them of $q-1$ colors. Using Lemma 14, we therefore obtain that

$$\begin{aligned} \sum_{\gamma' \sim \gamma} e^{|\gamma'|} w_{\gamma'} &\leq |\gamma|(\Delta+1) \sum_{k \geq 1} (e\Delta)^{k-1} (q-1)^k e^k e^{-\tau k} = \frac{|\gamma|(\Delta+1)}{e\Delta} \sum_{k \geq 1} (e\Delta)^k (q-1)^k e^k e^{-\tau k} \\ &\leq \frac{|\gamma|(\Delta+1)}{e\Delta} \sum_{k \geq 1} e^{-3k} \leq |\gamma|, \end{aligned}$$

so the Kotecký-Preiss condition is satisfied.

The Kotecký-Preiss condition, in turn, implies the polymer mixing condition (2) with $\theta = 1/e$ since $e \cdot x \leq e^x$ for $x \geq 1$. For the same reason (since e^x gets much bigger than x), it is easy to see that the polymer mixing condition is weaker than the Kotecký-Preiss condition.

2.2 | The polymer Markov chain

For each $v \in V(G)$, let $\mathcal{A}(v) = \{\gamma \in C(G) : v \in \gamma\}$ denote the collection of all polymers containing v and let $a(v) = \sum_{\gamma \in \mathcal{A}(v)} w_{\gamma}$. By applying (2) to the smallest γ containing v we have $a(v) \leq \theta < 1$ for

all $v \in V(G)$. Define the probability distribution ν_v on $\mathcal{A}(v) \cup \{\emptyset\}$ by $\nu_v(\gamma) = w_\gamma$ for $\gamma \in \mathcal{A}(v)$ and $\nu_v(\emptyset) = 1 - a(v)$.

The polymer dynamics on Ω are defined by the following transition rule from a configuration Γ_t to a configuration Γ_{t+1} :

Polymer dynamics

1. Choose $v \in V(G)$ uniformly at random. Let $\gamma_v \in \Gamma_t \cap \mathcal{A}(v)$ if $\Gamma_t \cap \mathcal{A}(v) \neq \emptyset$ and let $\gamma_v = \emptyset$ otherwise. Note that γ_v is well defined since Γ_t can have at most one polymer containing v .
2. Mutually exclusively do the following:
 - With probability $\frac{1}{2}$, let $\Gamma_{t+1} = \Gamma_t \setminus \gamma_v$.
 - With probability $\frac{1}{2}$, sample γ from ν_v , set $\Gamma_{t+1} = \Gamma_t \cup \gamma$ if this is in Ω and set $\Gamma_{t+1} = \Gamma_t$ otherwise.

Note that the polymer dynamics are aperiodic, since there are self-loops, and irreducible since we can transition from any $\Gamma \in \Omega$ to any $\Gamma' \in \Omega$ (e.g., via the empty set). Since the polymer dynamics are finite, irreducible, and aperiodic, they are also ergodic. Next, we observe that the stationary distribution of the polymer dynamics is μ_G by checking detailed balance. Note that each transition of the dynamics changes a configuration Γ by at most one polymer γ ; let $\Gamma' = \Gamma \cup \gamma$. Then

$$\frac{\mu_G(\Gamma')}{\mu_G(\Gamma)} = \frac{\prod_{\gamma' \in \Gamma'} w_{\gamma'}}{\prod_{\gamma' \in \Gamma} w_{\gamma'}} = w_\gamma = \frac{\frac{|\gamma|}{n} \cdot \frac{1}{2} \cdot w_\gamma}{\frac{|\gamma|}{n} \cdot \frac{1}{2}} = \frac{P_{\Gamma \rightarrow \Gamma'}}{P_{\Gamma' \rightarrow \Gamma}},$$

where P is the transition matrix of the polymer dynamics, and so μ_G is the stationary distribution.

We now formally define the mixing time. If \mathcal{M} is an ergodic Markov chain with transition matrix P and stationary distribution ν then the mixing time of \mathcal{M} from a state x is given by

$$T_x(\varepsilon) = \min\{t > 0 \mid \text{for all } t' \geq t, \|P^{t'}(x, \cdot) - \nu(\cdot)\|_{TV} \leq \varepsilon\},$$

where $\|\nu' - \nu\|_{TV}$ denotes the total variation distance between distributions ν and ν' . The mixing time of \mathcal{M} is given by $T_{\text{mix}}(\varepsilon) = \max_x T_x(\varepsilon)$. We will write $T_{\text{mix}}(\mathcal{M}, \varepsilon)$ below if we need to emphasize which Markov chain we refer to.

2.3 | Proof of Theorems 2 and 5

Theorem 2. *Suppose that a polymer model $(C(\cdot), w, \mathcal{G})$ satisfies the polymer mixing condition (2). Then for each $G \in \mathcal{G}$ there is a Markov chain making single polymer updates with stationary distribution μ_G and mixing time $T_{\text{mix}}(\varepsilon) = O(n \log(n/\varepsilon))$.*

Proof. We will show that under condition (2) the mixing time of the polymer dynamics is $O(n \log(n/\varepsilon))$ by applying the path coupling technique. We define a metric $D(\cdot, \cdot)$ on Ω by setting $D(\Gamma, \Gamma') = 1$ if $\Gamma' = \Gamma \cup \{\gamma\}$ or $\Gamma = \Gamma' \cup \{\gamma\}$ for a polymer γ and extending this as a shortest path metric; that is, $D(\Gamma, \Gamma') = |\Gamma \triangle \Gamma'|$ for any $\Gamma, \Gamma' \in \Omega$ where \triangle denotes the symmetric difference of two sets.

Now suppose we couple two chains X_t and Y_t by attempting the same updates in both chains at each step. Suppose that $X_t = Y_t \cup \{\gamma\}$ for some polymer γ . With probability $\frac{|\gamma|}{n} \cdot \frac{1}{2}$ we pick $v \in \gamma$ and remove γ_v which yields $X_{t+1} = Y_{t+1} = X_t$. On the other hand, we may attempt to add a polymer $\gamma' \sim \gamma$

so that $Y_t \cup \{\gamma'\} \in \Omega$. That is, $X_{t+1} = X_t = Y_t \cup \{\gamma\}$ and $Y_{t+1} = Y_t \cup \{\gamma'\}$. This occurs with probability $\frac{|\gamma'|}{n} \cdot \frac{1}{2} \cdot w_{\gamma'}$ and in this case $D(X_{t+1}, Y_{t+1}) \leq 2$. Putting these together we can bound

$$\mathbb{E}[D(X_{t+1}, Y_{t+1})] \leq 1 + \frac{1}{2n} \left[-|\gamma| + \sum_{\gamma' \sim \gamma} |\gamma'| w_{\gamma'} \right].$$

Using (2) we have $\sum_{\gamma' \sim \gamma} |\gamma'| w_{\gamma'} \leq \theta |\gamma|$ and so

$$\mathbb{E}[D(X_{t+1}, Y_{t+1})] \leq 1 - |\gamma| \frac{1 - \theta}{2n} \leq 1 - \frac{1 - \theta}{2n}.$$

By the path coupling lemma (see [12], Section 6), and with W denoting the diameter of Ω under $D(\cdot, \cdot)$, we have that the mixing time is at most $\log(W/\varepsilon)2n/(1 - \theta) = O(n \log(n/\varepsilon))$, using that $W \leq 2n$. This finishes the proof. ■

To prove Theorem 5 we will show that a single update of the polymer dynamics can be computed in constant expected time. Assume that the polymer model is computationally feasible and that the polymer sampling condition (3) holds with constant $\tau \geq 5 + 3 \log((q - 1)\Delta)$. We will use the following algorithm. Let $r = \tau - 2 - \log((q - 1)\Delta) \geq 3 + 2 \log((q - 1)\Delta)$ and let $\mathcal{A}_k(v) = \{\gamma \in \mathcal{A}(v) : |\gamma| \leq k\}$.

Single polymer sampler

1. Choose \mathbf{k} according to the following geometric distribution: for k a nonnegative integer,

$$\Pr[\mathbf{k} = k] = (1 - e^{-r})e^{-rk}.$$

This gives $\Pr[\mathbf{k} \geq k] = e^{-rk}$.

2. Enumerate all polymers in $\mathcal{A}_{\mathbf{k}}(v)$ and compute their weight functions.
3. Mutually exclusively output $\gamma \in \mathcal{A}_{\mathbf{k}}(v)$ with probability $w_{\gamma}e^{r|\gamma|}$, and with all remaining probability output \emptyset . In particular if $\mathbf{k} = 0$, then output \emptyset with probability 1.

In order to show that this algorithm has constant expected running time, we will require the following result on enumerating connected subgraphs of bounded degree graphs.

Lemma 15 ([28] Lemma 3.7). *Let G have maximum degree Δ and let $v \in V(G)$. There is an algorithm running in time $O(k^5(e\Delta)^{2k})$ that outputs a list of all connected subgraphs of G of size at most k containing v .*

We now proceed to prove the following lemma.

Lemma 16. *Under the polymer sampling condition (3) the output distribution of the single polymer sampler is ν_v . Furthermore, assuming the polymer model is computationally feasible, the expected running time of the sampler is constant.*

Proof. We first show that the probabilities $w_{\gamma}e^{r|\gamma|}$ sum to less than 1, which shows the last step of the sampling algorithm is well defined. Since $\tau - r = 2 + \log((q - 1)\Delta)$,

$$\sum_{\gamma \in \mathcal{A}(v)} w_{\gamma}e^{r|\gamma|} \leq \frac{1}{2} \sum_{k \geq 1} (e\Delta)^{k-1} (q - 1)^k e^{-\tau k + rk} = \frac{1}{2e\Delta} \sum_{k \geq 1} e^{-k} < 1.$$

We next show that the output of the algorithm has distribution ν_v . Given $\gamma \in \mathcal{A}(v)$, to output γ we must choose $\mathbf{k} \geq |\gamma|$. This happens with probability $e^{-r|\gamma|}$ by the distribution of \mathbf{k} . Conditioned on

choosing such a \mathbf{k} , the probability we output γ is $w_\gamma e^{r|\gamma|}$, and multiplying these probabilities together gives w_γ as desired. Since this is true for all $\gamma \in \mathcal{A}(v)$, the output distribution is exactly v_v .

Finally we analyze the expected running time assuming that the model is computationally feasible. To do this, we observe that by Lemma 15, conditioned on the event that $\mathbf{k} = k$ the enumeration step of the algorithm takes time $O(k^5(e\Delta)^{2k})$, and the time to determine which polymers are allowed and computing their weights is $O(k^c(q-1)^k(e\Delta)^{k-1}/2)$ for some $c > 0$, since the polymer model is computationally feasible; here, the factor k^c accounts for the time to determine whether a single polymer of size k is “allowed” and to compute its weight. Therefore, the expected running time is

$$\begin{aligned} &= O\left(1 + \sum_{k \geq 1} \Pr[\mathbf{k} = k] (k^5(e\Delta)^{2k} + k^c(e(q-1)\Delta)^k)\right) \\ &= O\left(1 + \sum_{k \geq 1} e^{-rk} k^c(e(q-1)\Delta)^{2k}\right) = O\left(1 + \sum_{k \geq 1} k^c e^{-(\tau'+1)k}\right) = O(1), \end{aligned}$$

where $\tau' = \tau - 5 - 3 \log((q-1)\Delta) \geq 0$. ■

Finally we prove Theorem 5.

Theorem 5. *If a computationally feasible polymer model $(C(\cdot), w, \mathcal{G})$ satisfies the polymer sampling condition (3) then for all $G \in \mathcal{G}$ there is an ε -approximate sampling algorithm for μ_G with running time $O(n \log(n/\varepsilon))$.*

Proof. By Theorem 2, there is $T_\varepsilon = O(n \log(n/\varepsilon))$ so that if we start with the empty configuration $\Gamma_0 = \emptyset$ and run the polymer dynamics, then Γ_{T_ε} has distribution within $\varepsilon/2$ total variation distance of μ_G .

By Lemma 16, in expectation the running time will be $O(n \log(n/\varepsilon))$, but we want an upper bound on the worst case running time as well. To do this, we will simply stop the algorithm and output the empty configuration if the total running time exceeds $L = Cn \log(n/\varepsilon)$, for some sufficiently large constant C . We next show that the probability that the algorithm does not terminate in L steps is at most $\varepsilon/2$, which therefore yields that the output distribution has total variation distance at most ε from μ_G .

The randomness in the running time comes from the choice of the geometric random variable \mathbf{k} at each step and the time taken to enumerate polymers in $\mathcal{A}_{\mathbf{k}}(v)$. By the choice of r , the random variable that takes the value $k^5(e\Delta)^{2k} + k^c(e(q-1)\Delta)^k$ with probability $(1 - e^{-r})e^{-rk}$ has exponential tails, and so a Chernoff bound (see, e.g., [20, Corollary 2.3]) shows that the probability that the sum of $\Theta(n \log(n/\varepsilon))$ independent copies of such a random variable is at least twice its expectation is bounded by $e^{-\Theta(n \log(n/\varepsilon))}$ which is at most $\varepsilon/2$ (for large enough choice of constants), finishing the proof. ■

3 | APPROXIMATE COUNTING ALGORITHM

In this section we show how to use a sampling oracle to approximately compute the partition function of the polymer model. One standard way is by self-reducibility. In [18] an efficient sampling algorithm for polymer models is derived from an efficient approximate counting algorithm by applying self-reducibility on the level of polymers. While we could apply polymer self-reducibility in the other direction to obtain counting algorithms from our sampling algorithm, here we use the simulated annealing method instead (see [2, 19, 31]) to obtain a faster implementation of counting from sampling.

Suppose that $(C(G), w)$ is a computationally feasible polymer model. Let ρ be a parameter and define a weight function

$$w_\gamma(\rho) = w_\gamma e^{-\rho|\gamma|}$$

for all $\gamma \in C(G)$. Then for each $\rho \geq 0$ this defines a computationally feasible polymer model $(C(G), w(\rho))$ on G , where setting $\rho = 0$ recovers the original model $(C(G), w)$. If the original model $(C(G), w)$ satisfies the polymer sampling condition (3), then so does $(C(G), w(\rho))$ for every $\rho \geq 0$ as the weight function $w_\gamma(\rho)$ is monotone decreasing in ρ .

Given the graph G , we write the partition function of the polymer model $(C(G), w(\rho))$ as a function of ρ :

$$Z(\rho) = Z(G; \rho) = \sum_{\Gamma \in \Omega} \prod_{\gamma \in \Gamma} w_\gamma(\rho) = \sum_{\Gamma \in \Omega} \prod_{\gamma \in \Gamma} w_\gamma e^{-\rho|\gamma|}.$$

The associated Gibbs distribution is denoted by $\mu_\rho = \mu_{G, \rho}$. Since $\lim_{\rho \rightarrow \infty} w_\gamma(\rho) = 0$, we have $\lim_{\rho \rightarrow \infty} Z(\rho) = 1$ (only the empty configuration Γ contributes to this limit), and so we will use simulated annealing to interpolate between $Z(\infty) = 1$ and our goal $Z(0)$, assuming access to a sampling oracle for $(C(G), w(\rho))$ for all $\rho \geq 0$. To apply the simulated annealing method, roughly speaking, we find a sequence of parameters $0 = \rho_0 < \rho_1 < \dots < \rho_\ell < \infty$ called a *cooling schedule* where $\ell \in \mathbb{N}^+$, and then estimate $Z(0)$ using the telescoping product

$$\frac{1}{Z(0)} = \frac{1}{Z(\rho_0)} = \frac{Z(\rho_1)}{Z(\rho_0)} \frac{Z(\rho_2)}{Z(\rho_1)} \dots \frac{Z(\rho_\ell)}{Z(\rho_{\ell-1})} \frac{1}{Z(\rho_\ell)}.$$

To estimate each term $Z(\rho_{i+1})/Z(\rho_i)$, we define independent random variables

$$W_i = \prod_{\gamma \in \Gamma_i} \frac{w_\gamma(\rho_{i+1})}{w_\gamma(\rho_i)}, \quad \text{where } \Gamma_i \sim \mu_{\rho_i}.$$

It is straightforward to see that $\mathbb{E}[W_i] = Z(\rho_{i+1})/Z(\rho_i)$ (see Lemma 17). Using the sampling oracle for μ_{ρ_i} , we can sample W_i for all i , and by taking the product we get an estimate for $1/Z(0)$.

The key ingredient of simulated annealing is finding a good cooling schedule. There are non-adaptive schedules [2] that depend only on n , and adaptive schedules [19, 31] that also depend on the structure of $Z(\cdot)$. Usually the latter leads to faster algorithms than the former. In this paper we will use a simple nonadaptive schedule: $\rho_i = i/n$ for $i = 0, \dots, \ell$ where $\ell = O(n \log(n/\varepsilon))$. We will show that this cooling schedule already gives us a fast algorithm for the polymer model. The reason behind it is that the weight function $w_\gamma(\rho)$ decays exponentially fast, and so (see Lemma 18) the partition function $Z(\rho_\ell)$ is bounded by a constant when $\rho_\ell = O(\log n)$, leading to a short cooling schedule.

Our algorithm is as follows.

Polymer approximate counting algorithm

1. Let $\rho_i = i/n$ for $i = 0, 1, \dots, \ell$ where $\ell = \lceil n \log(4e(q-1)\Delta n/\varepsilon) \rceil$;
2. For $j = 1, \dots, m$ where $m = \lceil 64\varepsilon^{-2} \rceil$:

(a) For $0 \leq i \leq \ell - 1$:

- (i) Sample $\Gamma_i^{(j)}$ from μ_{ρ_i} ;
- (ii) Let $W_i^{(j)} = \prod_{\gamma \in \Gamma_i^{(j)}} e^{-|\gamma|/n}$;

(b) Let $W^{(j)} = \prod_{i=0}^{\ell-1} W_i^{(j)}$;

3. Let $\hat{W} = \frac{1}{m} \sum_{j=1}^m W^{(j)}$ and output $\hat{Z} = 1/\hat{W}$.

Before proving Theorem 6, we first present a few useful lemmas. We shall use $\rho_i = i/n$ for $0 \leq i \leq \ell$ as our cooling schedule and we further define $\rho_{\ell+1} = (\ell + 1)/n$ though it does not appear in the algorithm. For $0 \leq i \leq \ell - 1$ independently we define Γ_i to be a random sample from μ_{ρ_i} and $W_i = \prod_{\gamma \in \Gamma_i} e^{-|\gamma|/n}$. Finally, we let $W = \prod_{i=0}^{\ell-1} W_i$.

Lemma 17. For $0 \leq i \leq \ell - 1$,

$$\mathbb{E}[W_i] = \frac{Z(\rho_{i+1})}{Z(\rho_i)} \quad \text{and} \quad \mathbb{E}[W_i^2] = \frac{Z(\rho_{i+2})}{Z(\rho_i)}.$$

Therefore,

$$\mathbb{E}[W] = \frac{Z(\rho_\ell)}{Z(0)} \quad \text{and} \quad \mathbb{E}[W^2] = \frac{Z(\rho_\ell)Z(\rho_{\ell+1})}{Z(0)Z(\rho_1)}.$$

Proof. In the proof, we use $W_i(\Gamma_i)$ to denote $\prod_{\gamma \in \Gamma_i} \frac{w_\gamma(\rho_{i+1})}{w_\gamma(\rho_i)}$. We deduce from the definition of W_i that

$$\begin{aligned} \mathbb{E}[W_i] &= \sum_{\Gamma_i \in \Omega} \mu_{\rho_i}(\Gamma_i) W_i(\Gamma_i) = \frac{1}{Z(\rho_i)} \sum_{\Gamma_i \in \Omega} \prod_{\gamma \in \Gamma_i} w_\gamma e^{-i|\gamma|/n} \prod_{\gamma \in \Gamma_i} e^{-|\gamma|/n} \\ &= \frac{1}{Z(\rho_i)} \sum_{\Gamma_i \in \Omega} \prod_{\gamma \in \Gamma_i} w_\gamma e^{-(i+1)|\gamma|/n} = \frac{Z(\rho_{i+1})}{Z(\rho_i)} \end{aligned}$$

and that

$$\begin{aligned} \mathbb{E}[W_i^2] &= \sum_{\Gamma_i \in \Omega} \mu_{\rho_i}(\Gamma_i) W_i(\Gamma_i)^2 = \frac{1}{Z(\rho_i)} \sum_{\Gamma_i \in \Omega} \prod_{\gamma \in \Gamma_i} w_\gamma e^{-i|\gamma|/n} \prod_{\gamma \in \Gamma_i} e^{-2|\gamma|/n} \\ &= \frac{1}{Z(\rho_i)} \sum_{\Gamma_i \in \Omega} \prod_{\gamma \in \Gamma_i} w_\gamma e^{-(i+2)|\gamma|/n} = \frac{Z(\rho_{i+2})}{Z(\rho_i)}. \end{aligned}$$

Since $W_0, \dots, W_{\ell-1}$ are mutually independent, we obtain

$$\mathbb{E}[W] = \prod_{i=0}^{\ell-1} \mathbb{E}[W_i] = \prod_{i=0}^{\ell-1} \frac{Z(\rho_{i+1})}{Z(\rho_i)} = \frac{Z(\rho_\ell)}{Z(\rho_0)}$$

and

$$\mathbb{E}[W^2] = \prod_{i=0}^{\ell-1} \mathbb{E}[W_i^2] = \prod_{i=0}^{\ell-1} \frac{Z(\rho_{i+2})}{Z(\rho_i)} = \frac{Z(\rho_\ell)Z(\rho_{\ell+1})}{Z(\rho_0)Z(\rho_1)}.$$

Lemma 18. Suppose that $w_\gamma \leq 1$ for all $\gamma \in C(G)$. Then we have

$$1 \leq Z(\rho_\ell) \leq e^{\varepsilon/2}.$$

Proof. It is trivial that $Z(\rho_\ell) \geq 1$ since $\emptyset \in \Omega$ has weight 1. Meanwhile, we have the crude bound

$$Z(\rho_\ell) \leq \prod_{\gamma \in \mathcal{C}(G)} (1 + w_\gamma e^{-\ell|\gamma|/n}).$$

We then deduce that

$$\begin{aligned} \log(Z(\rho_\ell)) &\leq \sum_{\gamma \in \mathcal{C}(G)} w_\gamma e^{-\ell|\gamma|/n} \leq \sum_{v \in V} \sum_{k \geq 1} \sum_{\substack{\gamma \in \mathcal{C}(G): \\ |\gamma|=k, v \in \gamma}} (e^{-\ell/n})^k \\ &\stackrel{(a)}{\leq} n \sum_{k \geq 1} \left(\frac{e(q-1)\Delta}{e^\ell/n} \right)^k \stackrel{(b)}{\leq} n \sum_{k \geq 1} \left(\frac{\varepsilon}{4n} \right)^k \leq n \cdot \frac{2\varepsilon}{4n} = \frac{\varepsilon}{2}, \end{aligned}$$

where (a) follows from Lemma 14 and (b) from $\ell \geq n \log(4e(q-1)\Delta n/\varepsilon)$. \blacksquare

Lemma 19. *We have*

$$\frac{Z(\rho_1)}{Z(0)} \geq \frac{1}{e} \quad \text{and} \quad \frac{Z(\rho_{\ell+1})}{Z(\rho_\ell)} \leq 1.$$

Proof. Since the weight function $w_\gamma(\rho)$ is decreasing in ρ , the partition function $Z(\rho)$ is also decreasing, which implies $Z(\rho_{\ell+1}) \leq Z(\rho_\ell)$. On the other hand, recalling Lemma 17, we have

$$\frac{Z(\rho_1)}{Z(0)} = \mathbb{E}[W_0] = \mathbb{E} \left[\prod_{\gamma \in \Gamma_0} e^{-|\gamma|/n} \right],$$

where Γ_0 is sampled from μ_{ρ_0} . Notice that for any $\Gamma_0 \in \Omega$ we have

$$\prod_{\gamma \in \Gamma_0} e^{-|\gamma|/n} = \exp \left(-\frac{1}{n} \sum_{\gamma \in \Gamma_0} |\gamma| \right) \geq \frac{1}{e}.$$

Thus, the lemma follows. \blacksquare

We are now ready to prove Theorem 6 which we restate for convenience.

Theorem 6. *If a computationally feasible polymer model $(C(\cdot), w, \mathcal{G})$ satisfies the polymer sampling condition (3) then for all $G \in \mathcal{G}$ there is a randomized ε -approximate counting algorithm for $Z(G)$ with running time $O((n/\varepsilon)^2 \log^2(n/\varepsilon))$ and success probability at least $3/4$.*

Proof. We first assume that we have access to an exact sampler $\mathcal{S}_{\text{exact}}$ that samples from μ_ρ for all $\rho \geq 0$. Using this sampler in the Polymer approximate counting algorithm, we find that, for each j and each i , $\Gamma_i^{(j)}$ is an exact sample from the distribution μ_{ρ_i} and hence $W_i^{(j)}$ is an exact sample of W_i , independently for every j and i . Thus, $W^{(j)}$ is a sample of W independently for every j , and \hat{W} is the sample mean of $W^{(j)}$'s. We deduce from Lemmas 17 and 18 that

$$(1 + \varepsilon/2)\mathbb{E}[W] \leq \frac{e^{\varepsilon/2}Z(\rho_\ell)}{Z(0)} \leq \frac{e^\varepsilon}{Z(0)},$$

and

$$(1 - \varepsilon/2)\mathbb{E}[W] \geq \frac{e^{-\varepsilon}Z(\rho_\ell)}{Z(0)} \geq \frac{e^{-\varepsilon}}{Z(0)},$$

where we use $1 + \varepsilon/2 \leq e^{\varepsilon/2}$ and $e^{-\varepsilon} \leq 1 - \varepsilon/2$ for all $0 < \varepsilon < 1$. Then

$$\Pr\left(\frac{e^{-\varepsilon}}{Z(0)} \leq \hat{W} \leq \frac{e^\varepsilon}{Z(0)}\right) \geq \Pr\left(\left|\hat{W} - \mathbb{E}[W]\right| \leq (\varepsilon/2)\mathbb{E}[W]\right).$$

By Chebyshev's inequality we have

$$\Pr\left(\left|\hat{W} - \mathbb{E}[W]\right| \geq (\varepsilon/2)\mathbb{E}[W]\right) \leq \frac{4 \operatorname{Var}(W)}{\varepsilon^2 m (\mathbb{E}[W])^2} \leq \frac{4(e-1)}{\varepsilon^2 m} \leq \frac{1}{8},$$

where the second to last inequality follows from Lemmas 17 and 19:

$$\frac{\operatorname{Var}(W)}{(\mathbb{E}[W])^2} = \frac{\mathbb{E}[W^2]}{(\mathbb{E}[W])^2} - 1 = \frac{Z(0)}{Z(\rho_1)} \frac{Z(\rho_{\ell+1})}{Z(\rho_\ell)} - 1 \leq e - 1.$$

Thus, we deduce that

$$\Pr\left(e^{-\varepsilon}Z(0) \leq \hat{Z} \leq e^\varepsilon Z(0)\right) = \Pr\left(\frac{e^{-\varepsilon}}{Z(0)} \leq \hat{W} \leq \frac{e^\varepsilon}{Z(0)}\right) \geq \frac{7}{8},$$

so the error probability is at most $1/8$. Note that the number of samples that we used is ℓm .

Now we replace the exact sampling oracle S_{exact} by an approximate one. For every $\rho \geq 0$, the polymer model $(C(G), w(\rho))$ is computationally feasible and satisfies the polymer sampling condition (3). Thus, for any $\rho \geq 0$, Theorem 5 gives a randomized algorithm S that outputs a $1/(8\ell m)$ -approximate sample from μ_ρ . We then couple S and S_{exact} optimally and run the algorithm with both S and S_{exact} simultaneously, so that for any $\rho \geq 0$ samples from S and S_{exact} for μ_ρ coincide with probability at least $1 - 1/(8\ell m)$. Let \mathcal{B} be the event that at least one of the ℓm samples from S in the algorithm does not couple with that from S_{exact} . Then a union bound yields $\Pr(\mathcal{B}) \leq 1/8$. Let \mathcal{F} be the event that the algorithm using S_{exact} fails. From our argument before we see that $\Pr(\mathcal{F}) \leq 1/8$. Note that if neither of \mathcal{B} and \mathcal{F} happens, then the algorithm with S will output a desired estimate. Hence, we conclude from the union bound that the algorithm with S fails with probability at most

$$\Pr(\mathcal{B}) + \Pr(\mathcal{F}) \leq \frac{1}{8} + \frac{1}{8} = \frac{1}{4}.$$

Finally, we consider the running time of our algorithm. By Theorem 5 the running time of step 2(a)(i) is $O(n \log(8\ell mn)) = O(n \log(n/\varepsilon))$, and for step 2(a)(ii) the running time is $O(n)$. Thus, the running time of the algorithm is upper bounded by $\ell m \cdot O(n \log(n/\varepsilon)) = O((n/\varepsilon)^2 \log^2(n/\varepsilon))$. ■

4 | APPLICATIONS

Here we apply our results on subset polymer models to several approximate counting and sampling problems at low temperatures.

4.1 | Ferromagnetic Potts model

In this section, we prove Theorem 9 for the Potts model. Throughout this section, we will work under the assumptions/conditions of Theorem 9. That is, we fix a real number $\alpha > 0$, integers $q \geq 2$ and $\Delta \geq 3$ and a real number $\beta \geq \frac{5+3\log((q-1)\Delta)}{\alpha}$. We let \mathcal{G} be the class of α -expander graphs G with maximum degree at most Δ .

Consider the polymer model defined in Example 2 on an n -vertex graph $G \in \mathcal{G}$ with $M = n/2$ and ground state color $g \in [q]$. We will use $C^g = C^g(G)$ to denote the polymers and w_γ^g to denote the weight of a polymer $\gamma \in C^g$; recall that $w_\gamma^g = e^{-\beta B(\gamma)}$, where $B(\gamma)$ counts the number of external edges of γ plus the number of bichromatic internal edges. Let $Z^g(G)$ be the partition function of the polymer model $(C^g(G), w^g)$.

Lemma 20. *Under the conditions of Theorem 9, the polymer model $(C^g(\cdot), w^g, \mathcal{G})$ satisfies the polymer sampling condition (3) with $\tau = \alpha\beta$.*

Proof. Since every $G \in \mathcal{G}$ is an α -expander, for $\gamma \in C^g$ we have $B(\gamma) \geq \alpha|\gamma|$ and hence $w_\gamma^g \leq e^{-\tau|\gamma|}$. ■

The following lemma is from [22].

Lemma 21 ([22], Lemma 12). *For any n -vertex α -expander graph G and $\beta \geq 2\log(eq)/\alpha$, $qZ^g(G)$ is an e^{-n} -approximation of the Potts partition function $Z_{G,\beta}$.*

We are now ready to prove Theorem 9.

Theorem 9. *Suppose $q \geq 2$, $\Delta \geq 3$ are integers and $\alpha > 0$ is a real. Then for $\beta \geq \frac{5+3\log((q-1)\Delta)}{\alpha}$ and any $qe^{-n} \leq \varepsilon < 1$, there is an ε -approximate sampling algorithm for the q -state ferromagnetic Potts model with parameter β on all n -vertex α -expander graphs of maximum degree Δ with running time $O(n\log(n/\varepsilon))$. There is also an ε -approximate counting algorithm with running time $O((n/\varepsilon)^2 \log^2(n/\varepsilon))$ and success probability at least $3/4$.*

Proof. Let \mathcal{G} be the class of α -expander graphs of maximum degree at most Δ . Clearly, the polymer models $(C^g(\cdot), w^g, \mathcal{G})$ are computationally feasible. By Lemma 20, the models also satisfy the polymer sampling condition and therefore Theorems 5 and 6 apply. Consider any n -vertex graph $G \in \mathcal{G}$. Since $\beta \geq \frac{5+3\log((q-1)\Delta)}{\alpha} > \frac{2\log(eq)}{\alpha}$, Lemma 21 applies to G .

For the sampling algorithm, we pick a color $g \in [q]$ uniformly at random and generate an (ε/q) -approximate sample from the Gibbs measure associated to $Z^g(G)$ using the algorithm of Theorem 5, in time $O(n\log(n/\varepsilon))$. By Lemma 21, we conclude that the resulting output is an ε -approximate sample for the Potts model.

For the counting algorithm, we pick an arbitrary $g \in [q]$ and produce using the algorithm of Theorem 6 a number \hat{Z} in time $O((n/\varepsilon)^2 \log^2(n/\varepsilon))$, which is an $\varepsilon/(2q)$ -approximation to $Z^g(G)$ with probability $\geq 3/4$. By Lemma 21, we conclude that $q\hat{Z}$ is an ε -approximation for the partition function of the Potts model (with the same probability). ■

4.2 | Hard-core model

In this section, we prove Theorem 11 for the hard-core model.

Suppose $G = (V^0, V^1, E)$ is an n -vertex bipartite α -expander graph of maximum degree Δ . We will consider the hard-core model on G at sufficiently large fugacities λ . There are two relevant ground

states corresponding to the two parts of G , one is the independent set given by V^0 and the other is given by V^1 . We will capture deviations from the two ground states using the “even” and “odd” polymer models of Jenssen, Keevash, and Perkins [22]. We remark that similar models were considered independently by Liao, Lin, Lu, and Mao [25].

For $i \in \{0, 1\}$, we say a set $S \subseteq V^i$ is *small* if $|S| \leq |V^i|/2$. In particular, Definition 10 requires that small sets expand.

Following [22], we will define a polymer model $(C^i(G^2), w^i)$; note that the host graph² is G^2 , rather than G . The set $C^i = C^i(G^2)$ of allowed polymers consists of all small sets $\gamma \subseteq V^i$ which are connected subgraphs in G^2 . The set of spins is $\{0, 1\}$ and the ground state spin for a vertex v is $g_v = 0$ if $v \in V^i$, and $g_v = 1$ if $v \in V^{1-i}$; the spin assignment σ_γ for a polymer γ gives the spin 1 to each $v \in \gamma$. The weight of a polymer $\gamma \in C^i$ is defined as

$$w_\gamma^i = \frac{\lambda^{|\gamma|}}{(1 + \lambda)^{N_G(\gamma)}}, \quad (4)$$

where we recall that $N_G(\gamma)$ denotes the set of vertices in G which are adjacent to some vertex in γ . The key observation behind the definition of the weights is that for a set Γ of compatible polymers from C^i , the contribution to $Z_{G,\lambda}$ of all independent sets I with $I \cap V^i = \bigcup_{\gamma \in \Gamma} \gamma$ is exactly

$$(1 + \lambda)^{|V^{1-i}|} \prod_{\gamma \in \Gamma} w_\gamma^i,$$

see [[22], Proof of Lemma 19] for more details.

Let $Z^i(G)$ denote the partition function of the polymer model (C^i, w^i) (where two polymers are compatible if their distance in the host graph G^2 is at least 2). Using that G is an α -expander, we have the following lemma from [22].

Lemma 22 ([22], Lemma 19). *For any $\lambda \geq e^{11/\alpha}$ and any n -vertex graph $G = (V^0, V^1, E)$ which is a bipartite α -expander, the number*

$$(1 + \lambda)^{|V^1|} Z^0(G) + (1 + \lambda)^{|V^0|} Z^1(G)$$

is an e^{-n} -approximation of the hard-core partition function $Z_{G,\lambda}$.

In particular, [22, Lemma 17] shows that $(1 + \lambda)^{|V^1|} Z^0(G) + (1 + \lambda)^{|V^0|} Z^1(G)$ counts the contribution to $Z_{G,\lambda}$ of every independent set I of G , but some independent sets are double counted: those independent sets I for which the 2-connected components of $V^0 \cap I$ and $V^1 \cap I$ are all small. We call these independent sets *sparse*. The proof of [22, Lemma 19] shows that the relative contribution to $Z_{G,\lambda}$ of sparse independent sets is at most e^{-n} .

We are now ready to prove Theorem 11.

Theorem 11. *Suppose $\Delta \geq 3$ is an integer and $\alpha \in (0, 1)$ is a real. Then for any $\lambda \geq (3\Delta)^{6/\alpha}$ and $4e^{-n} \leq \varepsilon < 1$, there is an ε -approximate sampling algorithm for the hard-core model with parameter λ on all n -vertex bipartite α -expander graphs of maximum degree Δ . There is also an ε -approximate counting algorithm for the hard-core model with success probability at least $1 - \varepsilon$. Both algorithms run in time $O((n/\varepsilon)^2 \log^3(n/\varepsilon))$.*

² G^2 is the graph on vertex set $V(G)$, where two vertices are connected if their distance in G is at most 2.

Proof. First note that $\lambda \geq (3\Delta)^{6/\alpha} \geq 9^{6/\alpha} > e^{11/\alpha}$, so Lemma 22 applies. Let \mathcal{G} denote the set of host graphs G^2 corresponding to bipartite α -expanders G of maximum degree Δ . Noting that the polymer models $(C^i(\cdot), w^i, \mathcal{G})$ are computationally feasible, we verify the polymer sampling condition (3) for them. Fix arbitrary $i \in \{0, 1\}$. As in [21, Section 4.2], we have the bound

$$w_{\gamma}^i = \frac{\lambda^{|\gamma|}}{(1+\lambda)^{N_G(\gamma)}} \leq \frac{\lambda^{|\gamma|}}{(1+\lambda)^{(1+\alpha)|\gamma|}} \leq \lambda^{-\alpha|\gamma|},$$

so, using that $\lambda \geq (3\Delta)^{6/\alpha}$, we have that the models satisfy the polymer sampling condition with $\tau = \alpha \log \lambda \geq 6 \log(3\Delta) \geq 5 + 3 \log \Delta^2$. Therefore, we may also apply Theorems 5 and 6.

For the counting algorithm, we apply Theorem 6. Namely, by taking the median of $O(\log(1/\epsilon))$ trials, we can obtain \hat{Z}^0 and \hat{Z}^1 which are $(\epsilon/32)$ -approximations to $Z^0(G)$ and $Z^1(G)$, respectively, with probability at least $1 - \epsilon/32$. Let \mathcal{E} be the event that \hat{Z}^0 and \hat{Z}^1 are indeed $(\epsilon/32)$ -approximations to $Z^0(G)$ and $Z^1(G)$. Conditioned on \mathcal{E} , the number

$$\hat{Z} = (1+\lambda)^{|\hat{V}^1|} \hat{Z}^0 + (1+\lambda)^{|\hat{V}^0|} \hat{Z}^1$$

is an $(\epsilon/32)$ -approximation to the number $A = (1+\lambda)^{|\hat{V}^1|} Z^0(G) + (1+\lambda)^{|\hat{V}^0|} Z^1(G)$. By Lemma 22 and since $\epsilon \geq 4e^{-n}$, A is an $(\epsilon/4)$ -approximation to $Z_{G,\lambda}$ and hence \hat{Z} is an ϵ -approximation to $Z_{G,\lambda}$. Since \mathcal{E} occurs with probability at least $1 - \epsilon/16$, we obtain that \hat{Z} is the desired approximation for the counting algorithm.

For the sampling algorithm, let \mathbf{i} be the random variable which takes the value 0 with probability $\frac{(1+\lambda)^{|\hat{V}^1|} \hat{Z}^0}{\hat{Z}}$ and the value 1 otherwise, where $\hat{Z}^0, \hat{Z}^1, \hat{Z}$ are the quantities computed earlier. Then, use Theorem 5 to obtain an $(\epsilon/8)$ -approximate sample from the Gibbs distribution corresponding to the polymer model $(C^i(G), w^i)$, say $\hat{\Gamma}^i$. Obtain then an independent set \hat{I} by including into \hat{I} each $v \in V^{1-\mathbf{i}} \setminus N_G(\bigcup_{\gamma \in \hat{\Gamma}^{\mathbf{i}}} \gamma)$ with probability $\frac{\lambda}{1+\lambda}$ and each vertex in $\bigcup_{\gamma \in \hat{\Gamma}^{\mathbf{i}}} \gamma$ (with probability 1). We claim that the output distribution of \hat{I} is ϵ -close to the hard-core distribution $\mu_{G,\lambda}$.

To prove this, consider the random independent set I obtained by repeating the same steps above but using instead perfectly accurate computations, that is, pick $i = 0$ with probability $\frac{(1+\lambda)^{|\hat{V}^1|} Z^0(G)}{A}$ and the value 1 otherwise, then, sample (perfectly) Γ^i from the Gibbs distribution corresponding to the polymer model $(C^i(G), w^i)$, and then obtain the independent set I by including into I each $v \in V^{1-i} \setminus N_G(\bigcup_{\gamma \in \Gamma^i} \gamma)$ with probability $\frac{\lambda}{1+\lambda}$ and each vertex in $\bigcup_{\gamma \in \Gamma^i} \gamma$ (with probability 1). Then, if I is not sparse, I is generated with probability $\lambda^{|I|}/A$ (cf. the observation below (4)). On the other hand, if I is sparse, then I is generated with probability $2\lambda^{|I|}/A$. But by Lemma 22 and the remark following, the total variation distance between the distribution of I and the hard-core distribution $\mu_{G,\lambda}$ is bounded by the relative weight of the sparse independent sets, which, by Lemma 22, is at most $e^{-n} \leq \epsilon/4$.

We next observe that, conditioned on the event \mathcal{E} (i.e., that \hat{Z}^0 and \hat{Z}^1 are $(\epsilon/32)$ -approximations to $Z^0(G)$ and $Z^1(G)$), there is a coupling between \hat{I} and I such that $\hat{I} = I$ with probability at least $1 - \epsilon/4$. Indeed, the total variation distance between \mathbf{i} and i is at most $e^{\epsilon/16} - 1 \leq \epsilon/8$ and hence there is a coupling of \mathbf{i} with i so that $\mathbf{i} = i$ with probability at least $1 - \epsilon/8$. Analogously, there is a coupling of $\hat{\Gamma}^{\mathbf{i}}$ with Γ^i so that $\hat{\Gamma}^{\mathbf{i}} = \Gamma^i$ with probability at least $1 - \epsilon/8$. Since \mathcal{E} occurs with probability at least $1 - \epsilon/16$, it follows that the overall total variation distance between \hat{I} and I is at most $\epsilon/2$.

Hence, the output distribution of \hat{I} is ϵ -close to the hard-core distribution $\mu_{G,\lambda}$, finishing the proof of Theorem 11. \blacksquare

5 | COMPARISON TO SPIN GLAUBER DYNAMICS

In this section, we derive results for spin Glauber dynamics, restricted to appropriate sets in the state space, based on our results above (using fairly standard Markov chain comparison techniques). We start with the general framework of subset polymer models and obtain Theorem 25, which is then applied to the ferromagnetic Potts and hard-core models.

5.1 | Restricted Glauber dynamics for polymer models

Here, we define the restricted Glauber dynamics for subset polymer models, and show the upcoming Theorem 25 which bounds its mixing time under some appropriate conditions.

Consider a subset polymer model as in Section 1.1. There is a natural map $f : \Omega \rightarrow \{0, 1, \dots, q-1\}^{V(G)}$ between allowed polymer configurations and spin configurations, given by $f(\Gamma)_v = \sigma_\gamma(v)$ if $\gamma \in \Gamma$ and $v \in \gamma$ and $f(\Gamma)_v = g_v$ if $v \notin \cup_{\gamma \in \Gamma} \gamma$. Let $\Omega_{\text{spin}} = f(\Omega)$ be the spin configurations obtainable as images of the map f . It will be helpful to consider the inverse map f^{-1} and extend its domain to all $\sigma \in \{0, 1, \dots, q-1\}^{V(G)}$, so that $f^{-1}(\sigma)$ is the polymer configuration consisting of polymers that are connected components of vertices which do not receive their ground state spin; note that the range of the extended f^{-1} is not limited to Ω anymore.

Restricted Glauber dynamics is defined as follows, starting from $\Gamma_t \in \Omega$.

1. Choose $v \in V(G)$ and $s \in \{0, \dots, q-1\}$ uniformly.
2. Γ' is formed from Γ_t by assigning v to spin s (formally, by letting $\sigma = f(\Gamma_t)$, forming σ' from σ by assigning v to spin s , and finally letting $\Gamma' = f^{-1}(\sigma')$).
3. If $\Gamma' \in \Omega$ let $p = \min(1, w(\Gamma')/w(\Gamma))$.
 - With probability p , $\Gamma_{t+1} = \Gamma'$.
 - With probability $1 - p$, $\Gamma_{t+1} = \Gamma_t$.
4. If $\Gamma' \notin \Omega$ then $\Gamma_{t+1} = \Gamma_t$.

We will use the Markov chain comparison technique to show that the restricted Glauber dynamics is rapidly mixing. To do this, we need a mild condition on the set of allowed polymers $\mathcal{C}(G)$. A polymer model is said to be *single-update-compatible* if, for every size- k polymer $\gamma \in \mathcal{C}(G)$, there is an ordering v_1, \dots, v_k of the vertices in γ such that, for all $i \in [k]$, the set $S_i = \{v_1, \dots, v_i\}$ induces a connected subgraph of γ and we have that $(S_i, \sigma_\gamma|_{S_i})$ is a valid polymer itself, that is, $(S_i, \sigma_\gamma|_{S_i}) \in \mathcal{C}(G)$.

We will use the comparison method of Diaconis and Saloff-Coste [8, 9] as applied to mixing times by Randall and Tetali [30]. In order to avoid discussion of eigenvalues here, we use the version from Observation 13 of the survey paper [11]. We first show that the restricted Glauber dynamics is a reversible ergodic Markov chain with stationary distribution μ_G , which is easy to see from its definition.

Lemma 23. *Let G be a graph and let $(\mathcal{C}(G), w)$ be a single-update-compatible polymer model. The restricted Glauber dynamics is ergodic and reversible with stationary distribution μ_G .*

Proof. The restricted Glauber dynamics is aperiodic since we remain in the same state with positive probability after performing an update. It is irreducible since we can transition from any $\Gamma \in \Omega$ to any $\Gamma' \in \Omega$ by adding and removing vertices one-by-one. This shows that the restricted Glauber dynamics is ergodic. To show that it is reversible and has stationary distribution μ_G , we check detailed balance. Suppose $\Gamma, \Gamma' \in \Omega$ with $\Gamma \neq \Gamma'$ and $P(\Gamma, \Gamma') > 0$ where P is the transition matrix of the restricted

Glauber dynamics. Then,

$$\frac{P(\Gamma, \Gamma')}{P(\Gamma', \Gamma)} = \frac{\frac{1}{n} \cdot \frac{1}{q} \cdot \min\{1, w(\Gamma')/w(\Gamma)\}}{\frac{1}{n} \cdot \frac{1}{q} \cdot \min\{1, w(\Gamma)/w(\Gamma')\}} = \frac{w(\Gamma')}{w(\Gamma)} = \frac{\mu_G(\Gamma')}{\mu_G(\Gamma)}.$$

The lemma follows. \blacksquare

We next give some standard definitions that will be used in our comparison proof. Let \mathcal{M} denote the restricted Glauber dynamics and P be its transition matrix. Let \mathcal{M}' be the polymer dynamics and denote its transition matrix by P' . Define $E^*(\mathcal{M})$ to be the set of pairs of configurations (Γ, Γ') that can be achieved by one transition of the restricted Glauber dynamics; that is, $E^*(\mathcal{M}) = \{(\Gamma, \Gamma') \in \Omega^2 : P(\Gamma, \Gamma') > 0\}$. Similarly, define $E^*(\mathcal{M}') = \{(\Gamma, \Gamma') \in \Omega^2 : P'(\Gamma, \Gamma') > 0\}$ for the polymer dynamics.

For every $(\Gamma, \Gamma') \in E^*(\mathcal{M}')$, we define a path $\mathcal{P}_{\Gamma, \Gamma'}$ from Γ to Γ' to be a sequence of configurations such that every adjacent pair is a transition of the restricted Glauber dynamics; that is, every adjacent pair of configurations is in $E^*(\mathcal{M})$. For this, we assume that the polymer model is single-update-compatible (see Section 1.3). If $\Gamma = \Gamma'$, then the choice is easy—we let $\mathcal{P}_{\Gamma, \Gamma'} = (\Gamma, \Gamma')$. Suppose instead that $\Gamma' = \Gamma \cup \gamma$ for some polymer $\gamma \in \Omega$. Recall that there is a natural one-to-one mapping $f : \Omega \rightarrow \Omega_{\text{spin}}$ between the set of all (polymer) configurations Ω and the set of spin configurations Ω_{spin} . Let $\sigma = f(\Gamma)$ and $\sigma' = f(\Gamma')$ be the corresponding spin configurations. If γ has size k , let v_1, \dots, v_k be the ordering of vertices of γ from the definition of single-update-compatible so that, for all $i \in [k]$, the polymer induced by vertices v_1, \dots, v_i is in $\mathcal{C}(G)$. Let $(\sigma = \sigma_0, \sigma_1, \dots, \sigma_k = \sigma')$ be the sequence of spin configurations such that each σ_j is obtained from σ_{j-1} by changing the spin of v_j from $\sigma(v) = g_v$ to $\sigma'(v)$. The path $\mathcal{P}_{\Gamma, \Gamma'}$ is then defined to be $(f^{-1}(\sigma_0), \dots, f^{-1}(\sigma_k))$. If $\Gamma' = \Gamma \setminus \gamma$ for some $\gamma \in \Omega$, we can define the path $\mathcal{P}_{\Gamma, \Gamma'}$ in a similar manner. Note that in both cases the length of the path is $|\mathcal{P}_{\Gamma, \Gamma'}| = k = |\gamma|$.

For every $(\Gamma_0, \Gamma'_0) \in E^*(\mathcal{M})$, the *congestion* of the edge (Γ_0, Γ'_0) is defined to be

$$A(\Gamma_0, \Gamma'_0) = \frac{1}{\mu_G(\Gamma_0)P(\Gamma_0, \Gamma'_0)} \sum_{\substack{(\Gamma, \Gamma') \in E^*(\mathcal{M}'): \\ \mathcal{P}_{\Gamma, \Gamma'} \ni (\Gamma_0, \Gamma'_0)}} \mu_G(\Gamma)P'(\Gamma, \Gamma')|\mathcal{P}_{\Gamma, \Gamma'}|.$$

The *congestion* of the choice of paths is the quantity

$$A = \max_{(\Gamma_0, \Gamma'_0) \in E^*(\mathcal{M})} A(\Gamma_0, \Gamma'_0).$$

The following comparison lemma gives an upper bound on the mixing time of the restricted Glauber dynamics by the mixing time of the polymer dynamics.

Lemma 24 ([11], Observation 13). *Let $c_1 = \min_{\Gamma \in \Omega} P(\Gamma, \Gamma)$ and $c_2 = \min_{\Gamma \in \Omega} \mu_G(\Gamma)$. Then, for any $0 < \varepsilon < 1$ we have*

$$T_{\text{mix}}(\mathcal{M}, \varepsilon) \leq \max \left\{ A \left(T_{\text{mix}} \left(\mathcal{M}', \frac{1}{2e} \right) + 1 \right), \frac{1}{2c_1} \right\} \ln \frac{1}{\varepsilon c_2}.$$

We now proceed to establish the mixing-time of the restricted Glauber dynamics, which is the main result of this section. We will apply this to both the hard-core model (on bipartite α -expander

graphs) and the ferromagnetic Potts model (on α -expander graphs), for which we will define appropriate single-update compatible polymer models. Furthermore, in both of these applications, M below will be logarithmic in n/ε , giving polynomial mixing time for the restricted Glauber dynamics.

Theorem 25. *Suppose that a polymer model $(C(\cdot), w, \mathcal{G})$ satisfies the polymer mixing condition. Consider a graph $G \in \mathcal{G}$ such that $(C(G), w)$ is single-update-compatible. Let $M = \max\{|\gamma| : \gamma \in C(G)\}$. Suppose that, for every pair of configurations $\Gamma, \Gamma' \in \Omega$ whose corresponding spin configurations $f(\Gamma), f(\Gamma') \in \Omega_{\text{spin}}$ differ at exactly one vertex, we have*

$$\frac{1}{\eta} \leq \frac{\mu_G(\Gamma)}{\mu_G(\Gamma')} \leq \eta \quad (5)$$

for some constant $\eta > 1$. Then for any $0 < \varepsilon < 1$, the restricted Glauber dynamics has mixing time

$$T_{\text{mix}}(\varepsilon) \leq O\left(M\eta^{M+1}n^2 \log n \log(\eta/\varepsilon)\right).$$

Proof. By Lemma 24, it suffices to upper bound the congestion $A(\Gamma_0, \Gamma'_0)$ for every $(\Gamma_0, \Gamma'_0) \in E^*(\mathcal{M})$ where

$$A(\Gamma_0, \Gamma'_0) = \sum_{\substack{(\Gamma, \Gamma') \in E^*(\mathcal{M}'): \\ \mathcal{P}_{\Gamma, \Gamma'} \ni (\Gamma_0, \Gamma'_0)}} \frac{\mu_G(\Gamma)}{\mu_G(\Gamma_0)} \cdot \frac{P'(\Gamma, \Gamma')}{P(\Gamma_0, \Gamma'_0)} \cdot |\mathcal{P}_{\Gamma, \Gamma'}|.$$

If $\Gamma_0 = \Gamma'_0$, then for our choices of paths to get $(\Gamma_0, \Gamma'_0) \in \mathcal{P}_{\Gamma, \Gamma'}$ we must have $\Gamma = \Gamma' = \Gamma_0 = \Gamma'_0$. It follows that

$$A(\Gamma_0, \Gamma_0) = \frac{P'(\Gamma_0, \Gamma_0)}{P(\Gamma_0, \Gamma_0)} \leq \frac{1}{1/q} = q$$

since $P(\Gamma_0, \Gamma_0) \geq 1/q$ by the update rule of the restricted Glauber dynamics.

Now suppose $\Gamma_0 \neq \Gamma'_0$. Let $\sigma_0 = f(\Gamma_0)$ and $\sigma'_0 = f(\Gamma'_0)$ be the corresponding spin configurations. Notice that σ_0 and σ'_0 differ at exactly one vertex, which we denote by v . If $\sigma_0(v) \neq g_v$ and $\sigma'_0(v) \neq g_v$ then no path $\mathcal{P}_{\Gamma, \Gamma'}$ would contain (Γ_0, Γ'_0) by our choice of paths, and thus $A(\Gamma_0, \Gamma'_0) = 0$. Assume next that $\sigma_0(v) = g_v$ and $\sigma'_0(v) \neq g_v$. Then, if $(\Gamma_0, \Gamma'_0) \in \mathcal{P}_{\Gamma, \Gamma'}$ for some $(\Gamma, \Gamma') \in E^*(\mathcal{M}')$, we must have $\Gamma' = \Gamma \cup \gamma$ for some polymer $\gamma \in \Omega$ and also $v \in \gamma$. Moreover, the spin configurations $f(\Gamma), f(\Gamma'), \sigma_0, \sigma'_0$ are all the same outside γ . This implies that the number of such paths $\mathcal{P}_{\Gamma, \Gamma'}$ is upper bounded by the number of polymers containing v .

Now fix some $(\Gamma, \Gamma') \in E^*(\mathcal{M}')$ such that $(\Gamma_0, \Gamma'_0) \in \mathcal{P}_{\Gamma, \Gamma'}$ and assume that $\Gamma' = \Gamma \cup \gamma$ for some polymer $\gamma \in \Omega$ with $v \in \gamma$. Then,

$$|\mathcal{P}_{\Gamma, \Gamma'}| \leq |\gamma| \leq M. \quad (6)$$

As the path $\mathcal{P}_{\Gamma, \Gamma'}$ is obtained by changing the spins vertex by vertex in the corresponding spin configurations, $f(\Gamma)$ and $f(\Gamma_0)$ differ at most $|\gamma|$ vertices. The condition of the theorem implies that

$$\frac{\mu_G(\Gamma)}{\mu_G(\Gamma_0)} \leq \eta^{|\gamma|} \leq \eta^M. \quad (7)$$

The update rule of the restricted Glauber dynamics gives

$$P(\Gamma_0, \Gamma'_0) = \frac{1}{n} \cdot \frac{1}{q} \cdot \min \left\{ 1, \frac{w(\Gamma'_0)}{w(\Gamma_0)} \right\} \geq \frac{1}{\eta q n} \quad (8)$$

and for the polymer dynamics we have

$$P'(\Gamma, \Gamma') = \frac{|\gamma|}{n} \cdot \frac{1}{2} \cdot w_\gamma = \frac{|\gamma|w_\gamma}{2n}. \quad (9)$$

Let γ_v denote a polymer on $\{v\}$ with a spin from $\{0, \dots, q-1\} \setminus g_v$. Then, the polymer mixing condition implies that $\sum_{\gamma \sim \gamma_v} |\gamma|w_\gamma \leq \theta|\gamma_v| < 1$ for some $\theta \in (0, 1)$. Combining this and inequalities (6), (7), (8), and (9), we get

$$A(\Gamma_0, \Gamma'_0) \leq \sum_{\gamma: v \in \gamma} \eta^M \cdot \eta q n \cdot \frac{|\gamma|w_\gamma}{2n} \cdot M = \frac{1}{2} q M \eta^{M+1} \sum_{\gamma \sim \gamma_v} |\gamma|w_\gamma \leq \frac{1}{2} q M \eta^{M+1}.$$

For the case where $\sigma_0(v) \neq g_v$ and $\sigma'_0(v) = g_v$, the proof is almost the same and we can get the same bound. Thus,

$$A = \max_{(\Gamma_0, \Gamma'_0) \in E^*(\mathcal{M})} A(\Gamma_0, \Gamma'_0) \leq \max \left\{ q, \frac{1}{2} q M \eta^{M+1} \right\} \leq q M \eta^{M+1}.$$

The theorem then follows from Theorem 2 and Lemma 24 once we notice that $P(\Gamma, \Gamma) \geq 1/q$ and that $\mu_G(\Gamma) \geq 1/(\eta q)^n$ for all $\Gamma \in \Omega$. ■

5.2 | Truncated polymer model

The bound on the mixing time of the restricted Glauber dynamics in Theorem 25 is exponential in the size of the largest polymer which is in general undesirable. For example, in our applications in Section 4, M was linear in the number of vertices of the host graph. Here, we show that, under the polymer sampling condition, we can restrict our attention to polymers of size $O(\log n)$ in the sense that the partition function as well as the Gibbs distribution of the truncated polymer model are close to those of the original polymer model.

Let $(\mathcal{C}(G), w)$ be a polymer model on a graph G . For $k > 0$, define the truncated polymer model $(\mathcal{C}_k(G), w)$ by

$$\mathcal{C}_k(G) = \{\gamma \in \mathcal{C}(G) : |\gamma| \leq k\}.$$

Also we let

$$\Omega_k = \{\Gamma \in \Omega \mid \Gamma \subseteq \mathcal{C}_k(G)\} = \{\Gamma \subseteq \mathcal{C}_k(G) \mid \forall \gamma, \gamma' \in \Gamma, \gamma \sim \gamma'\}$$

be the set of allowed configurations (note that $\Omega_k \subseteq \Omega$). The partition function of the truncated polymer model $(\mathcal{C}_k(G), w)$ is given by

$$Z_k(G) = \sum_{\Gamma \in \Omega_k} \prod_{\gamma \in \Gamma} w_\gamma.$$

The corresponding Gibbs distribution on Ω_k is defined by $\mu_{G,k}(\Gamma) = \frac{\prod_{\gamma \in \Gamma} w_\gamma}{Z_k(G)}$. We remark that if the original polymer model satisfies the polymer sampling condition then so does the truncated polymer model, and thus Theorem 5 also applies to the truncated model.

The following lemma asserts that the Gibbs distribution and the partition function of the truncated polymer model $(\mathcal{C}_k(G), w)$ are close to those of the original model $(\mathcal{C}(G), w)$, provided that the polymer sampling condition (3) holds.

Lemma 26. *Let \mathcal{G} be a family of graphs of maximum degree at most Δ and let $(C(\cdot), w, \mathcal{G})$ be a polymer model that satisfies the polymer sampling condition (3) with constant $\tau \geq 5 + 3 \log((q-1)\Delta)$. Let G be an n -vertex graph from \mathcal{G} . Then for any $\varepsilon > 0$ and $k = \frac{3 \log(2n/\varepsilon)}{2\tau}$, we have*

$$Z_k(G) \leq Z(G) \leq e^\varepsilon Z_k(G).$$

Moreover, the total variation distance between μ_G and $\mu_{G,k}$ is at most ε .

Proof. Note that $Z_k(G) \leq Z(G)$ follows immediately from $\Omega_k \subseteq \Omega$. For $\Gamma \in \Omega_k$, let $\Omega(\Gamma) = \{\Gamma' \in \Omega : \Gamma' \cap C_k(G) = \Gamma\}$ and let

$$\zeta(\Gamma) = \sum_{\Gamma' \in \Omega(\Gamma)} \prod_{\gamma \in \Gamma'} w_\gamma, \text{ so that } Z(G) = \sum_{\Gamma \in \Omega_k} \zeta(\Gamma). \quad (10)$$

Let $C_k^+(G) = C(G) \setminus C_k(G) = \{\gamma \in C(G) : |\gamma| > k\}$ be the collection of all polymers of size greater than k . Notice that for each $\Gamma \in \Omega_k$ we have the crude bound

$$\zeta(\Gamma) \leq \prod_{\gamma \in \Gamma} w_\gamma \prod_{\gamma \in C_k^+(G)} (1 + w_\gamma). \quad (11)$$

Combining (10) and (11), we obtain that

$$Z(G) \leq \prod_{\gamma \in C_k^+(G)} (1 + w_\gamma) \sum_{\Gamma \in \Omega_k} \prod_{\gamma \in \Gamma} w_\gamma = Z_k(G) \prod_{\gamma \in C_k^+(G)} (1 + w_\gamma). \quad (12)$$

Since $(1+x) \leq e^x$ for all real x , we have that

$$\log \left(\prod_{\gamma \in C_k^+(G)} (1 + w_\gamma) \right) \leq \sum_{\gamma \in C_k^+(G)} w_\gamma \leq \sum_{\gamma \in C_k^+(G)} e^{-\tau|\gamma|}. \quad (13)$$

The last inequality follows from the fact that $(C(\cdot), w, \mathcal{G})$ satisfies the polymer sampling condition with constant τ . Then we deduce from Lemma 14 that

$$\sum_{\gamma \in C_k^+(G)} e^{-\tau|\gamma|} = \sum_{v \in V} \sum_{\ell \geq k} \sum_{\substack{\gamma \in C_k^+(G): \\ |\gamma| = \ell, v \in \gamma}} e^{-\tau\ell} \leq n \sum_{\ell \geq k} \left(\frac{e(q-1)\Delta}{e^\tau} \right)^\ell, \quad (14)$$

and since $\tau \geq 5 + 3 \log((q-1)\Delta)$, we get $e(q-1)\Delta \leq e^{\tau/3}$. It follows that

$$\sum_{\ell \geq k} \left(\frac{e(q-1)\Delta}{e^\tau} \right)^\ell \leq 2 \exp \left(-\frac{2}{3} \tau k \right) = \frac{\varepsilon}{n}. \quad (15)$$

Combining (12), (13), (14), and (15) yields $Z(G) \leq e^\varepsilon Z_k(G)$, as needed. Finally, we bound the total variation distance between μ_G and $\mu_{G,k}$:

$$\|\mu_G - \mu_{G,k}\|_{TV} = \mu_G(\Omega \setminus \Omega_k) = \frac{Z(G) - Z_k(G)}{Z(G)} \leq 1 - e^{-\varepsilon} \leq \varepsilon,$$

where the first equality is because $\mu_G(\Gamma) > \mu_{G,k}(\Gamma)$ if and only if $\Gamma \in \Omega \setminus \Omega_k$, for which we have $\mu_{G,k}(\Gamma) = 0$. This finishes the proof. \blacksquare

5.3 | Applications

In this section, we apply the previous results to show that (spin) Glauber dynamics for the ferromagnetic Potts and hard-core models mix in polynomial time on expander graphs, when restricted to configurations close to the ground states (which, as we have already seen, constitute the main portion of the probability space at low temperatures).

5.3.1 | Restricted Glauber for ferromagnetic Potts

In this section, we prove Theorem 13 for the q -color ferromagnetic Potts model. Throughout this section, we will work under the assumptions/conditions of Theorem 9. That is, we fix a real number $\alpha > 0$, integers $q \geq 2$ and $\Delta \geq 3$ and a real number $\beta \geq \frac{5+3\log((q-1)\Delta)}{\alpha}$. We let \mathcal{G} be the class of α -expander graphs G with maximum degree at most Δ .

Let G be an n -vertex graph in \mathcal{G} and let ε be a value in $(qe^{-n}, 1)$. As in Section 4.1, we will consider the polymer model (C^g, w^g) whose polymers are connected subgraphs of G with at most $n/2$ vertices, which are labeled by the remaining colors $[q] \setminus \{g\}$. In fact, following Section 5.2, we will work with a truncation of this model. Namely, for $M > 0$, let (C_M^g, w^g) be the polymer model on G restricted to polymers of size at most M .

Observation 27. For every $M > 0$, the set $\Omega_M^g(G)$, as defined in Section 1.3, is precisely the set of allowable polymer configurations in the truncated polymer model (C_M^g, w^g) .

Theorem 13. Suppose $q \geq 2$, $\Delta \geq 3$ are integers and $\alpha > 0$ is a real. Let $\beta \geq \frac{5+3\log((q-1)\Delta)}{\alpha}$ be a real number and $g \in [q]$. Then, for any n -vertex α -expander graph G of maximum degree Δ and any $\varepsilon \in (0, 1)$, for $M = O(\log(n/\varepsilon))$ the Glauber dynamics restricted to $\Omega_M^g(G)$ has mixing time $T_{\text{mix}}(\varepsilon)$ polynomial in n and $1/\varepsilon$.

Proof. We let \mathcal{G} be the class of α -expanders with maximum degree at most Δ . For the given n -vertex graph $G \in \mathcal{G}$, let $\mu_{G,M}^g$ be the Gibbs distribution of the polymer model $(C_M^g(G), w^g)$.

By Lemma 20, we have that $(C^g(\cdot), w^g, \mathcal{G})$ satisfies the polymer sampling condition with $\tau = \alpha\beta$ and hence so does the truncated polymer model $(C_M^g(\cdot), w^g, \mathcal{G})$. The result therefore follows by applying Theorem 25, after observing that (i) the polymer model $(C_M^g(G), w^g)$ is single-update-compatible (use DFS ordering), and (ii) for a pair of polymer configurations $\Gamma, \Gamma' \in \Omega_M^g$ whose corresponding spin configurations σ, σ' differ at a vertex, we have

$$\frac{\mu_{G,M}^g(\Gamma)}{\mu_{G,M}^g(\Gamma')} = \frac{e^{-\beta m(G, \sigma)}}{e^{-\beta m(G, \sigma')}} \in [1/\eta, \eta],$$

where $\eta = \exp(\beta\Delta)$ (since G has maximum degree Δ , changing the spin of a vertex can create at most Δ new monochromatic/bichromatic edges). This finishes the proof. ■

The following lemma justifies that the set $\Omega_M^g(G)$ with $M = O(\log(n/\varepsilon))$ constitutes for all but ε of the aggregate weight of colorings in the Potts distribution on G .

Lemma 28. Suppose $q \geq 2$, $\Delta \geq 3$ are integers and $\alpha > 0$ is a real. Let $\beta \geq \frac{5+3\log((q-1)\Delta)}{\alpha}$ be a real number and $g \in [q]$. Then, for any n -vertex α -expander graph G of maximum degree Δ and any $qe^{-n} \leq \varepsilon < 1$, for $M = \frac{3\log(4n/\varepsilon)}{2\alpha\beta}$, $qZ_M^g(G)$ is an ε -approximation of the Potts partition function $Z_{G,\beta}$.

Proof. Let $Z^g(G)$ be the partition function of the polymer model $(C^g(G), w^g)$. Since $\beta \geq \frac{5+3\log((q-1)\Delta)}{\alpha} > \frac{2\log(eq)}{\alpha}$ and $\varepsilon \geq qe^{-n}$, by Lemma 21 we have that $qZ^g(G)$ is an $(\varepsilon/2)$ -approximation to $Z_{G,\beta}$. If \mathcal{G} is the class of α -expanders with maximum degree at most Δ then by Lemma 20, we have that $(C^g(\cdot), w^g, \mathcal{G})$ satisfies the polymer sampling condition with $\tau = \alpha\beta$ and hence so does the truncated polymer model $(C_M^g(\cdot), w^g, \mathcal{G})$. It follows by Lemma 26 that, for $M = \frac{3\log(4n/\varepsilon)}{2\alpha\beta}$, $qZ_M^g(G)$ is an $(\varepsilon/2)$ -approximation to $qZ^g(G)$. Therefore, $qZ_M^g(G)$ is an ε -approximation to $Z_{G,\beta}$. ■

5.3.2 | Restricted Glauber dynamics for hard-core mixes in polynomial time

In this section, we state and prove the analogue of Theorem 13 for the hard-core model. In particular, let $G = (V^0, V^1, E)$ be an n -vertex α -expanding bipartite graph of maximum degree Δ , and for $i \in \{0, 1\}$ and $M > 0$, let $\Omega_M^i(G)$ denote the independent sets I whose deviations from the ground state V^i consists of small connected components, more precisely, $(V^i \setminus I) \cup (I \cap V^{1-i})$ consists of connected components of size at most M . Using similar methods to Section 5.3.1, we will show the following.

Theorem 29. Fix $\alpha > 0$, $\Delta \geq 3$ and $\lambda \geq (6\Delta)^{3+6/\alpha}$. For any n -vertex bipartite α -expander with maximum degree at most Δ and any $\varepsilon \in (0, 1)$ and $i \in \{0, 1\}$, with $M = O(\log(n/\varepsilon))$, the Glauber dynamics restricted to $\Omega_M^i(G)$ has mixing time $T_{\text{mix}}(\varepsilon)$ polynomial in n and $1/\varepsilon$.

As we shall see soon in the upcoming Lemma 32, and for λ large enough, the set $\Omega_M^0(G) \cup \Omega_M^1(G)$ for $M = \Theta(\log(n/\varepsilon))$ captures all but ε weight of the hard-core partition function and hence Theorem 29 can be used to obtain another polynomial time algorithm for the hard-core model on expanding graphs G in that regime.

To prove Theorem 29, it will be simpler to work with somewhat different polymer models than those in Section 4.2. These models were originally used in [21] (the conference version of [22]). For $i \in \{0, 1\}$, and following [21], we will define a polymer model $(C^i(G), w^i)$. The host graph will be G and the model will capture deviations from the ground state V^i : a polymer γ will be a connected set of vertices in G such that $(V^i \setminus I) \cup (I \cap V^{1-i}) = \gamma$ for some independent set I . Specifically, the set $C^i = C^i(G)$ of allowed polymers consists of all connected sets of vertices γ (in G) such that $|\gamma \cap V^i| \leq |V^i|/4$ and for any $v \in \gamma \cap V^{1-i}$, all of the neighbors of v (in G) are also in γ . The set of spins is $\{0, 1\}$ and the ground state spin for a vertex v is $g_v = 1$ if $v \in V^i$, and $g_v = 0$ if $v \in V^{1-i}$; the spin assignment σ_γ for a polymer γ is given by $1 - g_v$ for $v \in \gamma$. The weight of a polymer $\gamma \in C^i$ is defined as

$$w_\gamma^i = \frac{\lambda^{|\gamma \cap V^{1-i}|}}{\lambda^{|\gamma \cap V^i|}}.$$

The main observation behind the definition of the weight w_γ^i is that the weight of an independent set I such that $(V^i \setminus I) \cup (I \cap V^{1-i}) = \gamma$ is $\lambda^{|V^i|} w_\gamma^i$.

Following again Section 5.2, it will be relevant to consider, for $M > 0$, the truncated polymer model $(C_M^i(G), w^i)$ whose polymers are of size at most M ; observe that the set Ω_M^i defined above is precisely the set of allowable polymer configurations in the truncated polymer model. We next verify the polymer sampling condition (3) for these models and conclude the proof of Theorem 29.

Lemma 30. Fix $\alpha > 0$ and $\Delta \geq 3$. Let \mathcal{G} be the class of bipartite α -expanders with maximum degree at most Δ . For $\lambda \geq (6\Delta)^{3+6/\alpha}$ and $i \in \{0, 1\}$, the polymer model $(C^i(\cdot), w^i, \mathcal{G})$ satisfies the polymer sampling condition (3) with $\tau = \frac{\alpha}{2+\alpha} \log \lambda$.

Proof. We have $\tau \geq 5 + 3 \log \Delta$, so it suffices to show that for $G \in \mathcal{G}$ and $\gamma \in (C^i(G), w^i)$ it holds that $w_\gamma^i \leq e^{-\tau|\gamma|}$. For $v \in \gamma \cap V^{1-i}$ we have that all of the neighbors of v are also in γ and hence, by the α -expansion of G , we have that $|\gamma \cap V^i| \geq (1 + \alpha)|\gamma \cap V^{1-i}|$. This gives $(2 + \alpha)|\gamma \cap V^i| \geq (1 + \alpha)|\gamma|$ and therefore

$$w_\gamma^i = \frac{\lambda^{|\gamma \cap V^{1-i}|}}{\lambda^{|\gamma \cap V^i|}} \leq \frac{\lambda^{\frac{1}{1+\alpha}|\gamma \cap V^i|}}{\lambda^{|\gamma \cap V^i|}} \leq \lambda^{-\frac{\alpha}{2+\alpha}|\gamma|} = e^{-\tau|\gamma|}. \quad \blacksquare$$

Proof of Theorem 29. Let \mathcal{G} be the class of bipartite α -expanders with maximum degree at most Δ . Consider an n -vertex graph $G \in \mathcal{G}$ and let $\mu_{G,M}^i$ be the Gibbs distribution of the model $(C_M^i(G), w^i)$.

By Lemma 30, we have that $(C^i(\cdot), w^i, \mathcal{G})$ satisfies the polymer sampling condition with $\tau = \frac{\alpha}{2+\alpha} \log \lambda$ and hence so does the truncated polymer model $(C_M^i(\cdot), w^i, \mathcal{G})$. The result therefore follows by applying Theorem 25, after observing that (i) the polymer model $(C_M^i(G), w^i)$ is single-update-compatible (use DFS ordering), and (ii) for a pair of polymer configurations $\Gamma, \Gamma' \in \Omega_M^i$ whose corresponding independent sets I, I' differ in at most one vertex, we have

$$\frac{\mu_{G,M}^i(\Gamma)}{\mu_{G,M}^i(\Gamma')} = \frac{\lambda^{|I|}}{\lambda^{|I'|}} \in [1/\lambda, \lambda]. \quad \blacksquare$$

Finally, we justify that, for large enough λ and $M = \Theta(\log(n/\epsilon))$, the aggregate weight of independent sets in $\Omega_M^0(G) \cup \Omega_M^1(G)$ captures all but ϵ fraction of the hard-core partition function $Z_{G,\lambda}$. Let $Z^i(G)$ denote the partition function of the polymer model (C^i, w^i) and $Z_M^i(G)$ denote the partition function of the polymer model (C_M^i, w_M^i) . We will need the following lemma from [21].

Lemma 31 ([21], Lemmas 4.1 and 4.2). *For $\lambda > \max \left\{ (2e)^{\frac{8n}{an_0}}, (2e)^{\frac{8n}{an_1}}, (2e)^{(40/\alpha)} \right\}$, the number $\lambda^{n_0} Z^0(G) + \lambda^{n_1} Z^1(G)$ is a $(2e^{-n})$ -approximation of the hard-core partition function $Z_{G,\lambda}$, where $n_i = |V^i|$ for $i \in \{0, 1\}$.*

Lemma 32. *Fix $\alpha \in (0, 1)$ and $\Delta \geq 3$. There exists a constant $C > 0$ such that for $\lambda > (6C\Delta)^{3+6/\alpha}$ the following holds for all n -vertex bipartite α -expander graphs $G = (V^0, V^1, E)$ of maximum degree at most Δ . For all $\epsilon \in (4e^{-n}, 1)$ and $M = \frac{3(2+\alpha)\log(4n/\epsilon)}{2\alpha \log \lambda}$, the number*

$$\hat{Z} = \lambda^{n_0} Z_M^0(G) + \lambda^{n_1} Z_M^1(G)$$

is an ϵ -approximation of the hard-core partition function $Z_{G,\lambda}$, where $n_i = |V^i|$ for $i \in \{0, 1\}$.

Proof. Let $n_i = |V^i|$ for $i \in \{0, 1\}$ and observe that $n/n_0, n/n_1 \leq 3$ (using that G is an α -expander for $\alpha \in (0, 1)$, see [21] for details). Therefore, by taking C large enough, we have that for all $\lambda > (6C\Delta)^{3+6/\alpha}$ both Lemmas 30 and 31 apply. Let $\epsilon \in (4e^{-n}, 1)$.

By Lemma 31, we have that $\lambda^{n_0} Z^0(G) + \lambda^{n_1} Z^1(G)$ is an $(\epsilon/2)$ -approximation to $Z_{G,\lambda}$. By Lemma 30, we have that, for $i \in \{0, 1\}$, $(C^i(\cdot), w^i, \mathcal{G})$ satisfies the polymer sampling condition with $\tau = \frac{\alpha}{2+\alpha} \log \lambda$ and hence so does the truncated polymer model $(C_M^i(\cdot), w^i, \mathcal{G})$. (Here, as usual, we take \mathcal{G} to be the class of bipartite α -expanders with maximum degree at most Δ .) It follows by Lemma 26 that, for $M = \frac{3(2+\alpha)\log(4n/\epsilon)}{2\alpha \log \lambda}$, $Z_M^i(G)$ is an $(\epsilon/2)$ -approximation to $Z^i(G)$. Therefore, \hat{Z} is an ϵ -approximation to $Z_{G,\lambda}$. \blacksquare

ACKNOWLEDGMENTS

The paper reflects only the authors' views and not the views of the ERC or the European Commission. The European Union is not liable for any use that may be made of the information contained therein.

REFERENCES

1. A. Barvinok, *Combinatorics and complexity of partition functions*. Algorithms and Combinatorics, Springer International Publishing, 2017.
2. I. Bezáková, D. Štefankovič, V. V. Vazirani, and E. Vigoda, *Accelerating simulated annealing for the permanent and combinatorial counting problems*, SIAM J. Comput. **37** (2008), 1429–1454.
3. M. Bordewich, C. Greenhill, and V. Patel, *Mixing of the Glauber dynamics for the ferromagnetic Potts model*, Random Structures Algorithms. **48** (2016), 21–52.
4. C. Borgs, *Absence of zeros for the chromatic polynomial on bounded degree graphs*, Combin. Probab. Comput. **15** (2006), 63–74.
5. C. Borgs, J. T. Chayes, A. Frieze, J. H. Kim, P. Tetali, E. Vigoda, and V. H. Vu, *Torpid mixing of some Monte Carlo Markov chain algorithms in statistical physics*. In Proceedings of the 40th Annual IEEE Symposium on Foundations of Computer Science (FOCS), 1999, pp. 218–229.
6. C. Borgs, J. T. Chayes, J. Kahn, and L. Lovász, *Left and right convergence of graphs with bounded degree*, Random Structures Algorithms. **42** (2013), 1–28.
7. C. Borgs and J. Z. Imbrie, *A unified approach to phase diagrams in field theory and statistical mechanics*, Commun. Math. Phys. **123** (1989), 305–328.
8. P. Diaconis and L. Saloff-Coste, *Comparison techniques for random walk on finite groups*, Ann. Probab. **21** (1993), 2131–2156.
9. P. Diaconis and L. Saloff-Coste, *Comparison theorems for reversible Markov chains*, Ann. Appl. Probab. **3** (1993), 696–730.
10. R. L. Dobrushin, *Estimates of semi-invariants for the Ising model at low temperatures*, Transl. Amer. Math. Soc. Ser. **2** (1996), 59–82.
11. M. E. Dyer, L. A. Goldberg, M. Jerrum, and R. Martin, *Markov chain comparison*, Probab. Surveys. **3** (2006), 89–111.
12. M. E. Dyer and C. S. Greenhill, *Random walks on combinatorial objects*. London Mathematical Society Lecture Note Series, Cambridge University Press, pp. 101–136.
13. M. E. Dyer and C. S. Greenhill, *On Markov chains for independent sets*, J. Algorithms. **35** (2000), 17–49.
14. C. Efthymiou, T. P. Hayes, D. Štefankovič, E. Vigoda, and Y. Yin, *Convergence of MCMC and loopy BP in the tree uniqueness region for the hard-core model*. Proceedings of the 57th Annual IEEE Symposium on Foundations of Computer Science (FOCS), 2016, pp. 704–713.
15. R. Fernández, P. A. Ferrari, and N. L. Garcia, *Loss network representation of Peierls contours*, Ann. Probab. (2001), 902–937.
16. D. Galvin and P. Tetali, *Slow mixing of Glauber dynamics for the hard-core model on regular bipartite graphs*, Random Structures Algorithms. **28** (2006), 427–443.
17. C. Gruber and H. Kunz, *General properties of polymer systems*, Commun. Math. Phys. **22** (1971), 133–161.
18. T. Helmuth, W. Perkins, and G. Regts, *Algorithmic Pirogov-Sinai theory*. Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, STOC 2019, 2019, pp. 1009–1020.
19. M. Huber, *Approximation algorithms for the normalizing constant of Gibbs distributions*, Ann. Appl. Probab. **25** (2015), 974–985.
20. S. Janson, T. Łuczak, and A. Ruciński, *Random graphs*. Wiley Series in Discrete Mathematics and Optimization, Wiley, 2011.
21. M. Jenssen, P. Keevash, and W. Perkins, *Algorithms for #BIS-hard problems on expander graphs*. Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms (SODA), 2019, pp. 2235–2247.
22. M. Jenssen, P. Keevash, and W. Perkins, *Algorithms for #BIS-hard problems on expander graphs*, 2019, arXiv preprint:1807.04804v2.
23. R. Kotecký and D. Preiss, *Cluster expansion for abstract polymer models*, Commun. Math. Phys. **103** (1986), 491–498.
24. L. Laanait, A. Messenger, S. Miracle-Solé, J. Ruiz, and S. Shlosman, *Interfaces in the Potts model I: Pirogov-Sinai theory of the Fortuin-Kasteleyn representation*, Commun. Math. Phys. **140** (1991), 81–91.
25. C. Liao, J. Lin, P. Lu, and Z. Mao, *Counting independent sets and colorings on random regular bipartite graphs*, 2019, arXiv preprint arXiv:1903.07531.
26. J. Liu and P. Lu, *FPTAS for #BIS with degree bounds on one side*. Proceedings of the 47th Annual ACM on Symposium on Theory of Computing (STOC), 2015, pp. 549–556.

27. E. Mossel, D. Weitz, and N. Wormald, *On the hardness of sampling independent sets beyond the tree threshold*, Probab. Theory Related Fields. **143** (2009), 401–439.
28. V. Patel and G. Regts, *Deterministic polynomial-time approximation algorithms for functions and graph polynomials*, SIAM J. Comput. **46** (2017), 1893–1919.
29. S. A. Pirogov and Y. G. Sinai, *Phase diagrams of classical lattice systems*, Teoret. Mat. Fiz. **25** (1975), 358–369.
30. D. Randall and P. Tetali, *Analyzing Glauber dynamics by comparison of Markov chains*, J. Math. Phys. **41** (2000), 1598–1615.
31. D. Štefankovič, S. Vempala, and E. Vigoda, *Adaptive simulated annealing: A near-optimal connection between sampling and counting*, J. ACM. **56** (2009), 18.
32. D. Weitz, *Counting independent sets up to the tree threshold*. Proceedings of the 38th Annual ACM Symposium on Theory of Computing (STOC), 2006, pp. 140–149.

How to cite this article: Chen Z, Galanis A, Goldberg LA, Perkins W, Stewart J, Vigoda E. Fast algorithms at low temperatures via Markov chains[†]. *Random Struct Alg.* 2020;1–28. <https://doi.org/10.1002/rsa.20968>