

Fast and Perfect Sampling of Subgraphs and Polymer Systems

Antonio Blanca ✉

Department of Computer Science and Engineering,
Pennsylvania State University, University Park, PA, USA

Sarah Cannon ✉

Department of Mathematical Sciences, Claremont McKenna College, CA, USA

Will Perkins ✉

Department of Computer Science, Georgia Institute of Technology, Atlanta, GA, USA

Abstract

We give an efficient perfect sampling algorithm for weighted, connected induced subgraphs (or *graphlets*) of rooted, bounded degree graphs. Our algorithm utilizes a vertex-percolation process with a carefully chosen rejection filter and works under a percolation subcriticality condition. We show that this condition is optimal in the sense that the task of (approximately) sampling weighted rooted graphlets becomes impossible in finite expected time for infinite graphs and intractable for finite graphs when the condition does not hold. We apply our sampling algorithm as a subroutine to give near linear-time perfect sampling algorithms for polymer models and weighted non-rooted graphlets in finite graphs, two widely studied yet very different problems. This new perfect sampling algorithm for polymer models gives improved sampling algorithms for spin systems at low temperatures on expander graphs and unbalanced bipartite graphs, among other applications.

2012 ACM Subject Classification Theory of computation → Randomness, geometry and discrete structures; Theory of computation → Design and analysis of algorithms; Theory of computation → Graph algorithms analysis; Theory of computation → Random walks and Markov chains

Keywords and phrases Random Sampling, perfect sampling, graphlets, polymer models, spin systems, percolation

Digital Object Identifier 10.4230/LIPIcs.APPROX/RANDOM.2022.4

Category RANDOM

Related Version *Full Version:* <https://arxiv.org/abs/2202.05907> [6]

Funding *Sarah Cannon:* Supported in part by NSF grants DMS-1803325 and CCF-2104795.

Will Perkins: Supported in part by NSF grant DMS-1847451.

Acknowledgements This work was carried out as part of the AIM SQuaRE workshop “Connections between computational and physical phase transitions.” We thank Tyler Helmuth, Alexandre Stauffer, and Izabella Stuhl for many helpful conversations.

1 Introduction

Sampling is a fundamental computational task: given a specification of a probability distribution on a (large) set of combinatorial objects, output a random object with the specified distribution or with a distribution close to the specified distribution. This task becomes challenging when the specification of the distribution is much more succinct than the set of objects, and one wants to sample using time and space commensurate with the specification. Fundamental examples include sampling from Markov random fields and probabilistic graphical models and sampling substructures of graphs. We will address both of these examples here and connect them in a new way.



© Antonio Blanca, Sarah Cannon, and Will Perkins;
licensed under Creative Commons License CC-BY 4.0

Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2022).

Editors: Amit Chakrabarti and Chaitanya Swamy; Article No. 4; pp. 4:1–4:18



Leibniz International Proceedings in Informatics

LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

We consider a natural sampling problem: given a bounded-degree graph G , sample a *graphlet* (a connected, vertex-induced subgraph) of G containing a fixed vertex r with probability proportional to an exponential in the size of the subgraph. That is, sample a graphlet S containing vertex r with probability proportional to $\lambda^{|S|}$, where $\lambda > 0$ is a distribution parameter and $|S|$ denotes the number of vertices in S . In this paper we are concerned with small values of λ , where the expected size of a sampled graphlet is much smaller than the size of the graph.

Sampling graphlets is an important task in data science, network analysis, bioinformatics, and sociology, as it allows us to gain information about massive graphs from small sections of it; see, e.g., [38, 28, 41, 4]. A number of variants of the problem have consequently been studied, including sampling graphlets of a given size uniformly at random or sampling weighted graphlets of all sizes [46, 5, 37, 15, 10, 11, 43, 1, 42, 45, 12, 9]. The variant we consider here, i.e., sample a graphlet S with probability proportional to $\lambda^{|S|}$, arises as a key subroutine in recent sampling algorithms for spin systems (hard-core model, Ising model, Potts model, etc.) in the regime of strong interactions via *polymer models* described below in Section 1.2; see [32, 13, 40, 25, 7, 36, 35, 14, 17].

One major limitation of previous sampling algorithms for graphlets and polymer models (those in, e.g., [32, 40, 45, 16, 24], among others) is the use of exhaustive enumeration of graphlets of a given size; this requires restrictive parameter regimes or large polynomial running times, with the maximum degree Δ of the graph appearing in the exponent of the polynomial. Here we design a fast perfect sampling algorithm for weighted graphlets based on a vertex percolation process combined with a rejection filter. This method bypasses the enumeration barrier and allows us to design perfect sampling algorithms for a number of applications, substantially improving upon existing algorithms in three ways: 1) our algorithms have considerably faster running times, with no dependence on Δ in the exponent; 2) our algorithms return perfect, rather than approximate, samples from the desired distributions; and 3) our algorithms are conceptually simple and practical to implement.

Our algorithm proceeds as follows. First, run a vertex percolation process on the graph G beginning at vertex r in a breadth-first search manner, repeatedly adding each adjacent vertex to the graphlet with a carefully-chosen probability p . Once the percolation process terminates, the graphlet is accepted as the random sample with a certain probability that depends on the graphlet and rejected otherwise; if the graphlet is rejected, the algorithm restarts another percolation process from r . Because of the careful way we choose the percolation and rejection probabilities, we can prove the final accepted sample is drawn exactly from the desired distribution and the expected running time is bounded by a constant that depends only on λ and the maximum degree Δ .

1.1 Sampling rooted graphlets

Our key contribution is a new algorithm for perfectly sampling weighted graphlets containing a given vertex r . Formally, let $G = (V, E)$ be a finite or infinite graph of maximum degree Δ . For $r \in V$, let $\mathcal{S}(G, r)$ be the set of all connected, vertex-induced subgraphs of G containing r . (The subgraph induced by $U \subseteq V$ has vertex set U and includes all the edges of G with both endpoints in U .) We call r the root of G and the elements of $\mathcal{S}(G, r)$ graphlets rooted at r . For $\lambda > 0$ define the probability distribution $\nu_{G,r,\lambda}$ on $\mathcal{S}(G, r)$ by

$$\nu_{G,r,\lambda}(S) = \frac{\lambda^{|S|}}{Z_{G,r,\lambda}}, \quad \text{where} \quad Z_{G,r,\lambda} = \sum_{\hat{S} \in \mathcal{S}(G,r)} \lambda^{|\hat{S}|}. \quad (1)$$

The distribution is well defined when the normalizing constant $Z_{G,r,\lambda}$, known as the partition function, is finite. This is the case for every graph of maximum degree Δ and every r when λ is below the critical threshold:

$$\lambda_*(\Delta) = \frac{(\Delta - 2)^{\Delta-2}}{(\Delta - 1)^{\Delta-1}}; \quad (2)$$

(see Lemma 3 below). We give an efficient perfect sampling algorithm for $\nu_{G,r,\lambda}$ for $\lambda < \lambda_*(\Delta)$.

► **Theorem 1.** *Fix $\Delta \geq 3$ and let $\lambda < \lambda_*(\Delta)$. There is a randomized algorithm that for any graph $G = (V, E)$ of maximum degree Δ and any $r \in V$ outputs a graphlet distributed according to $\nu_{G,r,\lambda}$ with expected running time bounded by a constant that depends only on Δ and λ .*

We assume a model of computation that allows for querying the adjacency list of a given vertex in a bounded degree graph in constant time, the standard model used in the study of sublinear algorithms [27]. We also assume access to a stream of perfectly random real numbers in $[0, 1]$. The model of computation is chosen for consistency; in particular, our methods extend to other models, only requiring to adjust the running time to account for any additional computational overhead.

Previous algorithms to generate ε -approximate samples from $\nu_{G,r,\lambda}$ (e.g., those in [32, 45, 16, 24]) exhaustively enumerate all graphlets of size $\leq k$, for some k that depends on the error parameter ε that describes how accurate the sample must be. This results in algorithms with $(1/\varepsilon)^{O(\log \Delta)}$ running times. Applications such as sampling from polymer models require multiple samples from $\nu_{G,r,\lambda}$ and, consequently, have small error tolerance per sample; in particular, they require $\varepsilon \ll 1/n$, which results in inefficient algorithms with overall running time $n^{O(\log \Delta)}$. The algorithm in Theorem 1, on the other hand, is an exact sampler whose expected running time depends only on Δ and λ and thus provides a significant advantage in applications as we detail below.

We also show that Theorem 1 is sharp in two ways. First, we establish that there is no polynomial-time approximate sampling algorithm for $\nu_{G,r,\lambda}$ when $\lambda > \lambda_*(\Delta)$ for the class of graphs of maximum degree at most Δ unless $\text{RP}=\text{NP}$. Second, in the infinite setting, the normalizing constant $Z_{G,r,\lambda}$ may diverge (and consequently the distribution $\nu_{G,r,\lambda}$ is not well-defined) when $\lambda > \lambda_*(\Delta)$; conversely, we prove that $Z_{G,r,\lambda}$ is finite on every graph of maximum degree Δ when $\lambda \leq \lambda_*(\Delta)$.

► **Lemma 2.** *If for every finite graph $G = (V, E)$ of maximum degree $\Delta \geq 3$ and every $r \in V$, there is a polynomial-time approximate sampler for $\nu_{G,r,\lambda}$ when $\lambda > \lambda_*(\Delta)$, then $\text{RP}=\text{NP}$.*

► **Lemma 3.** *The partition function $Z_{G,r,\lambda}$ is finite for every (possibly infinite) graph $G = (V, E)$ of maximum degree Δ and every $r \in V$ if and only if $\lambda \leq \lambda_*(\Delta)$.*

The proofs of these lemmas are omitted, but can be found in the full version of this paper [6].

Finally, we mention that the algorithmic result in Theorem 1 cannot be extended even to the case $\lambda = \lambda_*(\Delta)$: for the infinite Δ -regular tree, we can show that the expected size of a graphlet sampled from $\nu_{G,r,\lambda}$ is infinite when $\lambda = \lambda_*(\Delta)$, and so it is impossible to have sampling algorithms with finite expected running time. In summary, the algorithm in Theorem 1 for $\lambda < \lambda_*(\Delta)$, combined with the hardness/impossibility results in Lemmas 2 and 3 for $\lambda > \lambda_*(\Delta)$, provide a resolution to the computational problem of sampling from $\nu_{G,r,\lambda}$ on graphs of maximum degree at most Δ .

As mentioned, our sampling algorithm is based on exploring the connected component of r in a vertex-percolation process. We carefully choose a specific percolation parameter $p \in (0, 1)$ as a function of λ and Δ (see Lemma 9). We then perform breadth-first search (BFS) from r , labeling each new vertex encountered “active” with probability p and “inactive” with probability $1 - p$ independently over all vertices; we continue the BFS exploring only unexplored neighbors of active vertices. In this way we uncover the “active” component of r , call it γ . We then accept γ with a given probability depending on λ , Δ , $|\gamma|$ and $|\partial\gamma|$, where $\partial\gamma$ denotes the set of vertices outside of γ that are adjacent to γ . If we reject γ , we begin again with a new percolation process. We note that only when $\lambda < \lambda_*(\Delta)$ does there exist a suitable percolation probability p that results in a subcritical percolation process, so that the size of the active component has finite expectation and exponential tails. Sampling a graphlet by exploring a random component and performing a rejection step has been used in the past (most notably in the recent work of Bressan [9] to sample uniformly random graphlets of size k ; see also [2]). The weighted model we sample from is particularly well suited to this type of exploration algorithm because of the direct connection to a subcritical percolation process.

We prove a more general version of Theorem 1 in Section 2, allowing for vertex-labeled graphlets and modifications of the weights by multiplication by a non-negative function bounded by 1. These generalizations are needed for the application to polymer models in Section 1.2.

1.2 Sampling from polymer models

We use our algorithm for sampling weighted rooted graphlets to design fast and perfect samplers for polymer models. Polymer models are systems of interacting geometric objects representing defects from pure ground states (i.e., most likely configurations) in spin systems on graphs in classical statistical physics [29, 39, 22]. These geometric objects are most often represented by vertex-labeled graphlets from a given host graph. Recently, polymer models have found application as an algorithmic tool to sample from spin systems on various classes of graphs in strong interaction regimes; see, e.g., [32, 13, 40, 16, 31, 24, 25, 7, 36, 35, 14, 20, 17]. In these applications, the problem of sampling weighted vertex-labeled rooted graphlets emerged as a significant computational barrier.

We will work with *subset polymer models* in which all polymers are vertex-labeled graphlets from a host graph $G = (V, E)$. These models were defined in [29] and generalized in [39]. Such a polymer model consists of:

- A set $\mathcal{C} = \mathcal{C}(G)$ of polymers, each of which is a graphlet in G with the vertices of the graphlet labeled with colors from a set Σ of size q .
- Weights $w_\gamma \geq 0$ for each $\gamma \in \mathcal{C}$.
- An incompatibility relation \approx defined by connectivity. We say two polymers $\gamma, \gamma' \in \mathcal{C}$ are *incompatible* and write $\gamma \approx \gamma'$ if the union of their corresponding vertices induces a connected subgraph in G . Otherwise they are *compatible* and we write $\gamma \sim \gamma'$.

Let $\Omega(\mathcal{C})$ denote the set of all sets of pairwise compatible polymers from \mathcal{C} . The polymer model is the Gibbs distribution μ on $\Omega(\mathcal{C})$ defined by

$$\mu(X) = \frac{\prod_{\gamma \in X} w_\gamma}{Z(\mathcal{C})}, \quad \text{where} \quad Z(\mathcal{C}) = \sum_{X \in \Omega(\mathcal{C})} \prod_{\gamma \in X} w_\gamma$$

is the polymer model partition function. We say the weights of a polymer model are *computationally feasible* if w_γ can be computed in time polynomial in $|\gamma|$. The size $|\gamma|$ of a polymer γ is the number of vertices in the corresponding graphlet.

■ **Table 1** Comparison of conditions and running times of known polymer sampling algorithms.

Condition	Bound on exponential decay of weights	Running Time	Type of Sampler
Kotecký–Preiss [32, 35, 13]	$w_\gamma \leq (e^2 q \Delta)^{- \gamma }$	$n^{O(\log \Delta)}$	approximate
Polymer Sampling [16, 24]	$w_\gamma \leq (e^5 q^3 \Delta^3)^{- \gamma }$	$O(n \log n)$	approximate
Clique dynamics [23]	$w_\gamma \leq (eq\Delta)^{- \gamma }$	$n^{O(\log \Delta)}$	approximate
This work (Theorem 4)	$w_\gamma \leq (eq\Delta)^{- \gamma }$	$O(n \log n)$	perfect

We will assume without loss of generality that all vertex-labeled graphlets of G , including each individual vertex $v \in V$, are elements of \mathcal{C} , by setting $w_\gamma = 0$ when necessary. We let \mathcal{C}_v be all polymers containing vertex v .

Algorithms for sampling polymer models fall into two classes: those based on truncating the so-called cluster expansion of a polymer model to approximate a partition function and using self-reducibility to sample, and those based on Markov chains on the set of collections of compatible polymers. The cluster expansion approach, while giving polynomial-time algorithms, generally is relatively inefficient, with the degree of the polynomial bound on the running time growing with the degree of the underlying graph; e.g., running time $n^{O(\log \Delta)}$ in n -vertex graph of maximum degree Δ . The Markov chain approach in principle can be much faster (near linear time in the size of the graph) but runs into one hitch: a stricter condition on the parameters of the model is needed to perform one update step of the Markov chain (the “polymer sampling condition” in [16, 24]). We solve this problem by adapting our rooted graphlet sampler to sample polymers models, leading to a near linear-time perfect sampling algorithm for polymer models under the least restrictive conditions known (see Table 1).

► **Theorem 4.** *Consider a subset polymer model on a family of n -vertex graphs of maximum degree Δ with computationally feasible weights satisfying:*

$$w_\gamma \leq \lambda^{|\gamma|} \text{ for all } \gamma \in \mathcal{C} \text{ where } \lambda < \lambda_*(\Delta, q) := \frac{(\Delta - 2)^{\Delta-2}}{q(\Delta - 1)^{\Delta-1}}. \tag{3}$$

Suppose further that for all vertex v ,

$$\sum_{\gamma \not\ni v} |\gamma| w_\gamma < 1 + \sum_{\gamma \in \mathcal{C}_v} w_\gamma. \tag{4}$$

Then, there is a perfect sampling algorithm for μ with expected running time $O(n \log n)$.

The threshold defined in (3) is the generalization of the critical threshold for rooted graphlet sampling to the labeled case (taking $q = 1$ recovers the definition in (2)). Theorem 4 improves upon the known results for sampling from polymer models in two ways. For a very general class of polymer models, our algorithm simultaneously provides *perfect sampling* with *near-linear running time* under weak conditions on the polymer weights. We now review previous works to illustrate these improvements; see the accompanying Table 1.

A number of conditions on polymer weights have been used to provide efficient sampling algorithms. The first papers in this direction (including [32, 35, 13]) used the Kotecký–Preiss condition for convergence of the cluster expansion of the polymer model partition function [39]: $\sum_{\gamma' \ni \gamma} w_{\gamma'} e^{|\gamma'|} \leq |\gamma| \quad \forall \gamma \in \mathcal{C}$. This condition is typically verified by ensuring that:

$$\sum_{\gamma \ni v} w_\gamma e^{|\gamma|} \leq 1 \quad \forall v \in V. \tag{5}$$

Since the number of vertex-labeled rooted graphlets of size k in a maximum degree Δ graph grows roughly like $(eq\Delta)^{k-1}$ (see [8]), weights of polymers of size k must decay roughly like $(e^2q\Delta)^{-k}$ for the polymer model to satisfy (5), with the extra factor of e coming from the exponential in the left hand side of the condition (5).

The major downside to the algorithms based on the cluster expansion, i.e., those using (5) or the Kotecký–Preiss condition, is that the running times obtained are of the form $n^{O(\log \Delta)}$. Subsequent works, namely [16, 24], addressed this downside but at the cost of a significantly stricter condition on the polymer weights.

In [16], the authors devised a new Markov chain algorithm for sampling from polymer models. The condition on the polymer weights for rapid mixing of this chain is somewhat less restrictive than the Kotecký–Preiss condition; it is the Polymer Mixing condition:

$$\sum_{\gamma' \approx \gamma} |\gamma'| w_{\gamma'} \leq \theta |\gamma| \quad \forall \gamma \in \mathcal{C} \text{ for some } \theta \in (0, 1). \quad (6)$$

This requires weights of polymers of size k to decay like $(eq\Delta)^{-k}$, a savings of a factor e in the base of the exponent over (5). However, to implement a single step of this Markov chain in constant expected time, a stronger condition (the Polymer Sampling condition) was required:

$$w_\gamma \leq (e^5 \Delta^3 q^3)^{-|\gamma|}. \quad (7)$$

This is a significant loss of a factor $e^3 \Delta^2 q^2$ in the base of the exponent compared to (5), but the resulting sampling algorithm does run in near linear time.

In [23], the authors use a different Markov chain condition, the Clique Dynamics condition, similar to (6), which requires weights of polymers of size k to decay like $(eq\Delta)^{-k}$, saving the same factor e over (5). Their running times, though, are again of the form $n^{O(\log \Delta)}$ since implementing one step of their Markov chain involves enumerating rooted polymers of size $O(\log n)$.

Our results are a “best-of-both-worlds” for polymer sampling: under the conditions (3) and (4) that both require polymer weights to decay like $(eq\Delta)^{-k}$ – the precise conditions are similar to but slightly less restrictive than the polymer mixing condition (6) – we obtain a near linear time algorithm. Moreover, unlike any of the previous results, our algorithm is a perfect sampler.

To conclude this section, we comment briefly on the algorithm we design to sample from μ . Our starting point is the polymer dynamics Markov chain from [16]. We use it to implement a Coupling from the Past (CFTP) algorithm (see [44]). To do so efficiently (in terms of the number of steps of the Markov chain), we design a new “bounding Markov chain” for the polymer dynamics, a method pioneered in [33, 30], and to implement each step of the Markov chain efficiently, we turn to our sampler for weighted rooted graphlets from Theorem 1.

1.3 Applications to spin systems

Our new algorithm for sampling subset polymer models can be used as a subroutine in essentially all previous applications of polymer models for spin system sampling at low temperatures, including those in [35, 13, 40, 16, 31, 24, 25, 14, 20, 17]. This results in faster sampling algorithms under less restrictive conditions on model parameters in all those settings. As examples, we fleshed out here the details in two of these applications; more details are provided in the full version of this paper [6].

Hard-core model on bipartite graphs. The hard-core model on a graph G is the probability distribution μ_G^{hc} on $\mathcal{I}(G)$, the set of all independent sets of G , with

$$\mu_G^{hc}(I) = \frac{\lambda^{|I|}}{Z_G^{hc}(\lambda)}, \quad \text{where} \quad Z_G^{hc}(\lambda) = \sum_{I \in \mathcal{I}(G)} \lambda^{|I|}. \quad (8)$$

The complexity of approximate counting and sampling from μ_G^{hc} on bounded-degree graphs is well understood: there is a computational threshold at some $\lambda_c(\Delta)$, with efficient algorithms for $\lambda < \lambda_c(\Delta)$ [49, 3, 18, 19] and hardness above the threshold (no polynomial-time algorithms unless NP=RP) [47, 26, 48]. However, on bipartite graphs, the complexity of these problems is unresolved and is captured by the class #BIS (approximately counting independent sets in bipartite graphs) defined by Dyer, Goldberg, Greenhill, and Jerrum [21].

Theorem 4 implies the existence of a fast perfect sampling algorithm for the hard-core model in a certain class of bipartite graphs called *unbalanced bipartite graphs*, considered in [13, 23].

► **Corollary 5.** *There is a perfect sampling algorithm for μ_G^{hc} running in expected time $O(n \log n)$ for n -vertex bipartite graphs G with bipartition (L, R) , with maximum degree Δ_L in L , maximum degree Δ_R in R , and minimum degree δ_R in R if*

$$\lambda(1 + (1 + e)(\Delta_L - 1)\Delta_R) < (1 + \lambda)^{\delta_R/\Delta_L}. \quad (9)$$

Approximate sampling algorithms with large polynomial run times were previously given for this problem when $6\lambda\Delta_L\Delta_R < (1 + \lambda)^{\delta_R/\Delta_L}$ in [13] and when $3.3353\lambda\Delta_L\Delta_R < (1 + \lambda)^{\delta_R/\Delta_L}$ in [23]. Our result applies to a comparable parameter range: inequality (9) holds, for instance, when $(1 + e)\lambda\Delta_L\Delta_R < (1 + \lambda)^{\delta_R/\Delta_L}$, or when $3\lambda\Delta_L\Delta_R < (1 + \lambda)^{\delta_R/\Delta_L}$ and $\Delta_L < 6$. More importantly, our algorithm is the first to achieve perfect sampling and near-linear running time.

Potts model on expander graphs. The Q -color ferromagnetic Potts model on a graph $G = (V, E)$ is the probability distribution μ_G^{potts} on the set of Q -colorings of the vertices of G ; i.e., $\{1, \dots, Q\}^V$. Each Q -coloring σ is assigned probability $\mu_G^{potts}(\sigma) \propto e^{\beta m(G, \sigma)}$, where $m(G, \sigma)$ is the number of monochromatic edges of G under the coloring σ and $\beta > 0$ is a model parameter. When the parameter β is large, and G has some structure (e.g., G is an expander graph), typical configurations drawn from μ_G^{potts} are dominated by one of the Q colors; that is, there is phase coexistence in the model. This enables sampling using subset polymer models.

Recall that an n -vertex graph $G = (V, E)$ is an α -expander if for all subsets $S \subseteq V$ with $|S| \leq n/2$, the number of edges in E with exactly one endpoint in S is at least $\alpha|S|$.

► **Corollary 6.** *Consider the Q -color ferromagnetic Potts model on an α -expander n -vertex graph of maximum degree Δ . Suppose*

$$\beta \geq \frac{1 + \log\left(\frac{\Delta+1}{e\Delta} + 1\right) + \log((Q-1)\Delta)}{\alpha}. \quad (10)$$

Then there is a sampling algorithm with expected running time $O(n \log n)$ that outputs a sample σ with distribution $\hat{\mu}$ so that $\|\hat{\mu} - \mu_G^{potts}\|_{TV} \leq e^{-\Omega(n)}$.

Previously, [16] provided a ε -approximate sample for μ_G^{potts} in time $O(n \log(n/\varepsilon) \log(1/\varepsilon))$ whenever $\beta \geq \frac{5+3 \log((Q-1)\Delta)}{\alpha}$. Condition (10) holds when $\beta \geq \frac{1.2+\log((Q-1)\Delta)}{\alpha}$, so our algorithm applies to a larger range of parameters and removes the dependence on ε from the running time. We do not achieve perfect sampling in this application only because the subset polymer models used give approximations of μ_G^{potts} rather than describing μ_G^{potts} exactly.

1.4 Sampling unrooted graphlets in finite graphs

As another application of our algorithm for sampling weighted rooted graphlets, we consider next the problem of sampling weighted *unrooted* graphlets in a finite graph. Given a finite graph G , let $\mathcal{S}(G)$ be the set of all graphlets of G . Define the distribution $\nu_{G,\lambda}$ on $\mathcal{S}(G)$ by

$$\nu_{G,\lambda}(\gamma) = \frac{\lambda^{|\gamma|}}{Z_{G,\lambda}}, \quad \text{where } Z_{G,\lambda} = \sum_{\gamma \in \mathcal{S}(G)} \lambda^{|\gamma|}.$$

Read-McFarland and Štefankovič [45] gave a polynomial-time approximate sampling algorithm for $\nu_{G,\lambda}$ for the class of maximum-degree Δ graphs when $\lambda < \lambda_*(\Delta)$ and prove that there is no such algorithm for $\lambda \in (\lambda_*(\Delta), 1)$ unless $\text{NP}=\text{RP}$ ¹. We give a new algorithm for this problem, covering the entire $\lambda < \lambda_*(\Delta)$ regime, and improving on the result of [45] in two ways: (i) our running time is constant in expectation (with no dependence on n), while the running time of the ε -approximate sampler in [45] is $n \cdot (1/\varepsilon)^{O(\log \Delta)}$; and (ii) our algorithm outputs a perfect sample instead of an approximate one (and thus the running time has no dependence on any approximation parameter).

► **Theorem 7.** *Fix $\Delta \geq 3$ and let $\lambda < \lambda_*(\Delta)$. Then for the class of finite graphs of maximum degree Δ there is a randomized algorithm running in constant expected time that outputs a perfect sample from $\nu_{G,\lambda}$. The expected running time is bounded as a function of Δ and λ .*

The algorithm we use for this theorem is a modification of the one for sampling rooted graphlets. We pick a uniformly random $v \in V$, run the same BFS percolation exploration, and accept the connected component of v with an adjusted probability (to account for the fact that a graphlet can be generated from any of its vertices). The acceptance probability is bounded away from 0 and so the algorithm runs in constant expected time. As mentioned earlier, the ε -approximate sampling algorithm from [45] is based on the exhaustive enumeration of all subgraphs of size $\leq k$, for some k that depends on ε . Our new algorithm entirely bypasses this enumeration barrier.

2 Graphlet sampling: algorithms

In this section we present our efficient perfect sampling algorithm for weighted, vertex-labeled graphlets containing a fixed vertex r from a maximum degree Δ graph; in particular, in Section 2.1, we prove a generalized version of Theorem 1 from the introduction. We also provide in Section 2.2 our algorithm for sampling weighted graphlets (i.e., the unrooted, unlabeled case) and establish Theorem 7.

2.1 Sampling rooted vertex-labeled graphlets

Let $G = (V, E)$ be a (possibly infinite) graph of maximum degree Δ . For $U \subseteq V$, let $G[U]$ denote the corresponding vertex-induced subgraph of G ; specifically, $G[U] = (U, E(U))$, where $E(U) \subseteq E$ is the set of edges of G with both endpoints in U . A vertex-induced subgraph is a *graphlet* if it is connected. For $r \in V$, let $\mathcal{S}(G, r)$ be the set of all graphlets of G that contain vertex r . We call the graphlets in $\mathcal{S}(G, r)$ the graphlets rooted at r .

¹ In [45], the threshold is incorrectly stated as $\lambda < \lambda_*(\Delta + 1)$; this is due to a minor error interchanging the infinite Δ -regular tree with the infinite Δ -ary tree; with this small correction their analysis goes through with the bound on λ as stated here.

Let $\Sigma = \{1, \dots, q\}$ be a set of vertex labels or colors, and let $\mathcal{S}(G, r, q) = \bigcup_{S \in \mathcal{S}(G, r)} \Sigma^S$ be the set of all vertex-labeled graphlets rooted at r . Given a real parameter $\lambda > 0$, we assign to each rooted vertex-labeled graphlet $\gamma \in \mathcal{S}(G, r, q) \cup \{\emptyset\}$ with $|\gamma|$ vertices the weight $w_\gamma = \lambda^{|\gamma|} f(\gamma)$, where $f : \mathcal{S}(G, r, q) \cup \{\emptyset\} \rightarrow [0, 1]$. Note that $0 \leq w_\gamma \leq \lambda^{|\gamma|}$, which will be important for later analysis.

Define the probability distribution $\nu_{G, r, \lambda}$ on $\mathcal{S}(G, r, q) \cup \{\emptyset\}$ by setting

$$\nu_{G, r, \lambda}(\gamma) = \frac{w_\gamma}{Z(G, r, \lambda)}, \tag{11}$$

where $Z(G, r, \lambda) = \sum_{\gamma' \in \mathcal{S}(G, r, q) \cup \{\emptyset\}} w_{\gamma'}$. We assume that G, f, q and λ are such that $Z(G, r, \lambda)$ is finite, so that this distribution is well defined. When $q = 1$ and $f(\gamma) = \mathbb{1}(\gamma \neq \emptyset)$, $\nu_{G, r, \lambda}$ corresponds exactly to the distribution defined in (1) over the unlabeled graphlets of G rooted at r .

We consider the problem of sampling from $\nu_{G, r, \lambda}$; this more general version of the sampling problem is later used as a subroutine for sampling polymer systems in Section 3. Let

$$\lambda_*(\Delta, q) := \frac{(\Delta - 2)^{\Delta-2}}{q(\Delta - 1)^{\Delta-1}};$$

cf., (2). Our main algorithmic result for sampling colored rooted graphlets is the following.

► **Theorem 8.** *Suppose $\Delta \geq 3$, $\lambda > 0$, and $q \geq 1$ are such that $\lambda < \lambda_*(\Delta, q)$ and let $a > 0$ be a fixed constant. There is a randomized algorithm to exactly sample from $\nu_{G, r, \lambda}$ for graphs G of maximum degree Δ and functions $f : \mathcal{S}(G, r, q) \cup \{\emptyset\} \rightarrow [0, 1]$ where $f(\gamma)$ is computable in time $O(|\gamma|^a)$; this randomized algorithm has expected running time bounded by $C \cdot Z_{G, r, \lambda}^{-1}$, where $C > 0$ is a constant that depends only on q, λ, Δ and a .*

Theorem 1 from the introduction corresponds to the special case when $q = 1$ and $f(\gamma) = \mathbb{1}(\gamma \neq \emptyset)$ (in this case $Z_{G, r, \lambda} \geq \lambda$). Other mild assumptions on the function f , e.g., $f(\emptyset) = 1$ or $f(r) = 1$, ensure that $Z_{G, r, \lambda}$ is bounded away from 0 and, consequently, that the sampling algorithm in the theorem has constant expected running time.

As a warm-up, let us consider first our algorithm for sampling labeled rooted graphlets on a finite graph $G = (V, E)$ with $f = 1$, and purposely omit certain non-essential implementation details for clarity. First, we find $p \in (0, 1)$ such that $\frac{p}{q}(1 - p)^{\Delta-2} = \lambda$; this choice of p will be justified in what follows. The algorithm then repeats the following process until a vertex-labeled graphlet is accepted:

1. Each vertex of the graph is independently assigned with probability p a uniform random color from $\{1, \dots, q\}$, or it is marked as “not colored” with the probability $1 - p$.
2. Let γ be the vertex-labeled graphlet from $\mathcal{S}(G, r, q) \cup \{\emptyset\}$ corresponding the colored connected component of r ; i.e., the set of vertices connected to r by at least one path of colored vertices.
3. Observe that the probability of obtaining γ is $(p/q)^{|\gamma|}(1 - p)^{|\partial\gamma|}$, where $\partial\gamma$ denotes to set of vertices in G that are not in γ but are adjacent to a vertex in γ (with a slight abuse of notation, we let $|\partial\emptyset| = 1$). Our aim is to output γ with probability $\propto \lambda^{|\gamma|}$ which has no dependence on $\partial\gamma$. Therefore, we use a “rejection filter” and only accept γ with probability $(1 - p)^{(\Delta-2)|\gamma|+2-|\partial\gamma|}$, so that the probability that γ is the output becomes:

$$\left(\frac{p}{q}\right)^{|\gamma|} (1 - p)^{|\partial\gamma|} (1 - p)^{(\Delta-2)|\gamma|+2-|\partial\gamma|} = (1 - p)^2 \left(\frac{p}{q}(1 - p)^{\Delta-2}\right)^{|\gamma|} = (1 - p)^2 \lambda^{|\gamma|}. \tag{12}$$

From (12), the choice of p such that $\frac{p}{q}(1 - p)^{\Delta-2} = \lambda$ is apparent. We will prove that only when $\lambda < \lambda_*(\Delta, q)$ there exists $p \in (0, 1)$ such that $\frac{p}{q}(1 - p)^{\Delta-2} = \lambda$. In the actual implementation of the algorithm, it will in fact suffice to find an approximation for p .

We comment briefly on the intuition for the rejection filter. The acceptance probability must include a factor of $(1-p)^{-|\partial\gamma|}$, so that the final acceptance probability depends on $|\gamma|$ but not on $|\partial\gamma|$. However, $(1-p)^{-|\partial\gamma|} > 1$ is not a valid probability, so we use instead $(1-p)^{(\Delta-2)|\gamma|+2-|\partial\gamma|}$, which is at most 1 since $(\Delta-2)|\gamma|+2 \geq |\partial\gamma|$. This bound on $|\partial\gamma|$ is best possible since it is tight for the Δ -regular tree. We note that using looser bounds for $|\partial\gamma|$ affects the range of the parameter λ for which we can find $p \in (0, 1)$ so that $\frac{p}{q}(1-p)^{\Delta-2} = \lambda$.

Finally, we mention that the algorithm as described requires $\Omega(|V|)$ time per iteration and cannot be extended to infinite graphs. This is easily corrected by assigning colors starting from r and revealing only the colored component of r in a breadth-first fashion. The threshold $\lambda_*(\Delta, q)$ is sharp in the sense that only when $\lambda < \lambda_*(\Delta, q)$ is the value of p such that the revealing process is a sub-critical process that creates a small component with high probability. This ensures the algorithm can be implemented efficiently. In particular, we stress that our algorithm avoids exhaustively enumerating labeled graphlets, as done in previous methods [16].

Before giving the implementation details of our algorithm and proving Theorem 8, we consider the problem of finding $p \in (0, 1)$ such that $\frac{p}{q}(1-p)^{\Delta-2} = \lambda$. For $\Delta \geq 3$ and $q \geq 1$, consider the real function $g(x) = \frac{x}{q}(1-x)^{\Delta-2}$. It can be readily checked that the function g is continuous and differentiable in $[0, 1]$, has a unique maximum at $x = \frac{1}{\Delta-1}$ with $g(\frac{1}{\Delta-1}) = \lambda_*(\Delta, q)$, is increasing in $[0, \frac{1}{\Delta-1}]$, and decreasing in $[\frac{1}{\Delta-1}, 1]$. This implies that only when $\lambda < \lambda_*(\Delta, q)$, there exists a value of $p \in [0, \frac{1}{\Delta-1})$ such that $g(p) = \lambda$. In particular, when $\lambda > \lambda_*(\Delta, q)$, there is no value of p for which $g(p) = \lambda$ and when $\lambda = \lambda_*(\Delta, q)$, the only possible value is $p = \frac{1}{\Delta-1}$. The latter case would result in a *critical* percolation process, corresponding to the fact that the expected size of a graphlet from $\nu_{G,r,\lambda}$ has no uniform upper bound in the class of graphs of maximum degree Δ ; in fact, it is infinite on the Δ -regular tree. We can find a suitable approximation for p when $\lambda < \lambda_*(\Delta, q)$ via a simple (binary search) procedure.

► **Lemma 9.** *For any $\lambda \in [0, \lambda_*(\Delta, q))$ we can find rational numbers $\hat{p} \in [0, \frac{1}{\Delta-1})$ and $\hat{\lambda} \in [\lambda, \lambda_*(\Delta, q)]$ such that $g(\hat{p}) = \hat{\lambda}$ in $O(|\log \frac{1}{\Delta q(\lambda_* - \lambda)}|)$ time.*

The proof of this lemma appears after the proof of Theorem 8. We now prove Theorem 8, including giving a more detailed version of the algorithm outlined above that includes the previously omitted implementation details and allows for general functions $f : \mathcal{S}(G, r, q) \cup \{\emptyset\} \rightarrow [0, 1]$.

Proof of Theorem 8. For ease of notation, let $\lambda_* = \lambda_*(\Delta, q)$. Our algorithm to sample from $\nu_{G,r,\lambda}$ when $\lambda < \lambda_*$ explores from r in a breadth-first manner and stops once it has revealed the colored connected component of r . It proceeds as follows:

1. Find $\hat{p} \in [0, \frac{1}{\Delta-1})$ and $\hat{\lambda} \in [\lambda, \lambda_*)$ such that $g(\hat{p}) = \hat{\lambda}$. This can be done in time $O(|\log \frac{1}{\Delta q(\lambda_* - \lambda)}|)$ per Lemma 9.
2. Let Q be a queue. With probability $1 - \hat{p}$ do not add r to Q ; otherwise, assign r a color uniformly at random from $\{1, \dots, q\}$ and add r to Q . Mark r as explored.
3. While $Q \neq \emptyset$, repeat the following:
 - (3.1) Pop a vertex v from Q .
 - (3.2) For each unexplored neighbor w of v , with probability $1 - \hat{p}$ do not add w to Q ; otherwise, assign w a color uniformly at random from $\{1, \dots, q\}$ and add w to Q . Mark w as explored (regardless of whether it was added to Q or not).

4. Let γ be the vertex-labeled graphlet from $\mathcal{S}(G, r, q) \cup \{\emptyset\}$ corresponding the colored connected component of r . Accept γ with probability:

$$f(\gamma) \cdot (1 - \hat{p})^{(\Delta-2)|\gamma|+2-|\partial\gamma|} \left(\frac{\lambda}{\hat{\lambda}}\right)^{|\gamma|}.$$

5. If γ is rejected, go to Step 2 and repeat.

The probability of obtaining $\gamma \in \mathcal{S}(G, r, q) \cup \{\emptyset\}$ in an iteration of the algorithm is:

$$\left(\frac{\hat{p}}{q}\right)^{|\gamma|} (1 - \hat{p})^{|\partial\gamma|} \cdot f(\gamma) (1 - \hat{p})^{(\Delta-2)|\gamma|+2-|\partial\gamma|} \left(\frac{\lambda}{\hat{\lambda}}\right)^{|\gamma|} = (1 - \hat{p})^2 f(\gamma) \lambda^{|\gamma|} = (1 - \hat{p})^2 w_\gamma,$$

and thus the overall acceptance probability in an iteration is:

$$\rho := (1 - \hat{p})^2 \sum_{\gamma \in \mathcal{S}(G, r, q) \cup \{\emptyset\}} w_\gamma = (1 - \hat{p})^2 Z_{G, r, \lambda}.$$

Then,

$$\Pr[\gamma \in \mathcal{S}(G, r, q) \cup \{\emptyset\} \text{ is the output}] = \sum_{t \geq 1} (1 - \hat{p})^2 w_\gamma (1 - \rho)^{t-1} = \frac{(1 - \hat{p})^2 w_\gamma}{\rho} = \nu_{G, r, \lambda}(\gamma).$$

We next bound the expected running time of the algorithm. We claim first that expected running per iteration is at most a constant that depends only on a , Δ and q . If γ is the configuration generated in an iteration, it is discovered in $O(|\gamma| + |\partial\gamma|) = O(|\gamma|)$ time and, by assumption, $f(\gamma)$ can be computed in at most $O(|\gamma|^a)$ time, for suitable a constant $a > 0$. Let $\hat{\mu}$ the output distribution of Step 3 of the algorithm. Then, there exists a constant $C = C(q, \Delta) > 0$ such that the expected running time of each iteration is at most:

$$C \sum_{\gamma \in \mathcal{S}(G, r, q) \cup \emptyset} |\gamma|^{\max\{a, 1\}} \Pr_{\hat{\mu}}[\gamma] = C \cdot \mathbb{E}_{\hat{\mu}}[|\gamma|^{\max\{1, a\}}]. \quad (13)$$

We show next that $|\gamma|$ (under $\hat{\mu}$) is stochastically dominated by a random variable $W = X + Y$ (i.e., $|\gamma| \prec W$), where X and Y are i.i.d. random variables corresponding to the cluster size of a homogeneous Galton-Watson tree with offspring distribution $Bin(\Delta - 1, \hat{p})$. To see this, first note that $|\gamma| \prec L$, where L is the cluster size of a heterogeneous Galton-Watson tree, in which the root vertex has offspring distribution $Bin(\Delta, \hat{p})$ and every other vertex has offspring distribution $Bin(\Delta - 1, \hat{p})$. This is because the branching process generating γ includes the root only with probability \hat{p} (the root is always present in the Galton-Watson tree), and, in addition, it considers at most Δ (from the root) or $\Delta - 1$ (from any other vertex) potential branches. In turn, we have that $L \prec X + Y$, since we can couple the first $\Delta - 1$ branches of the root with X (starting at the root) and the remaining branch with Y (starting at the child of the root not coupled with X).

It is well-known that X and Y have finite moments when $(\Delta - 1)\hat{p} < 1$ (see, e.g., [34]). In particular, there exists a constant $A = A(a, \Delta, \hat{p}) > 0$ such that

$$\mathbb{E}_{\hat{\mu}}[|\gamma|^a] \leq \mathbb{E}[L^a] \leq \mathbb{E}[(X + Y)^a] \leq 2^a (\mathbb{E}[X^a] + \mathbb{E}[Y^a]) \leq A. \quad (14)$$

This together with (13) shows that the expected running time in each iteration of the algorithm is bounded by $C \cdot A$.

Now, let R be the number of times Steps 2–5 are repeated, let T be the overall the running time of the algorithm. Then:

$$\mathbb{E}[T] = \sum_{t \geq 1} \mathbb{E}[T \mid R = t] \Pr[R = t] \leq C \cdot A \cdot \sum_{t \geq 1} t(1 - \rho)^{t-1} \rho \leq \frac{CA}{\rho}, \quad (15)$$

and the result follows. ◀

We conclude this section with the proof of Lemma 9.

Proof of Lemma 9. It suffices to find $\hat{p} \in [p, \frac{1}{\Delta-1}]$. This can be done via binary search in t steps, provided $t \geq 0$ is such that $\frac{1}{\Delta-1} \cdot \frac{1}{2^t} \leq \frac{1}{\Delta-1} - p$. Since $g' \leq \frac{1}{q}$, it follows from the mean value theorem that $q(\lambda_* - \lambda) \leq \frac{1}{\Delta-1} - p$. Thus for the binary search to require at most t steps it is sufficient to pick t so that $\frac{1}{\Delta-1} \cdot \frac{1}{2^t} \leq q(\lambda_* - \lambda)$, and the result follows. \blacktriangleleft

2.2 Sampling unrooted graphlets

We consider next the problem of sampling weighted graphlets from a finite graph $G = (V, E)$ of maximum degree Δ ; specifically, in this variant of the sampling problem we consider unrooted, unlabeled, weighted graphlets of G . Let $\mathcal{S}(G)$ be the set of all graphlets of G . We define the probability distribution $\nu_{G,\lambda}$ on $\mathcal{S}(G)$ by setting

$$\nu_{G,\lambda}(S) = \frac{\lambda^{|S|}}{Z_{G,\lambda}},$$

where $Z_{G,\lambda} = \sum_{S' \in \mathcal{S}(G)} \lambda^{|S'|}$. The problem of (approximately) sampling from $\nu_{G,\lambda}$ is quite natural. In [45], it was established that this problem is computationally hard when $\lambda > \lambda_*(\Delta) = \frac{(\Delta-2)^{\Delta-2}}{(\Delta-1)^{\Delta-1}}$; an ε -approximate sampling algorithm was also given in [45] for the case when $\lambda < \lambda_*(\Delta)$ with running time $n \cdot (1/\varepsilon)^{O(\log \Delta)}$. We now provide the proof of Theorem 7, which says we can perfectly sample from $\nu_{G,\lambda}$ in constant expected time when $\lambda < \lambda_*(\Delta)$.

Proof of Theorem 7. For ease of notation, we set $\lambda_* = \lambda_*(\Delta)$ throughout this proof. Our algorithm to sample from $\nu_{G,\lambda}$ is based on the algorithm to sample from $\nu_{G,r,\lambda}$ (the rooted, vertex-labeled, weighted case). The idea is to pick a root uniformly at random and run the algorithm for the rooted case from this random vertex with the rejection filter adjusted to account for the fact that a graphlet can be generated from any of its vertices. It proceeds as follows:

1. Find $\hat{p} \in [0, \frac{1}{\Delta-1}]$ and $\hat{\lambda} \in [\lambda, \lambda_*)$ such that $g(\hat{p}) = \hat{\lambda}$ using the method from Lemma 9.
2. Pick a vertex $r \in V$ uniformly at random.
3. Let Q be a queue. With probability \hat{p} add r to Q and mark it as colored. Mark r as explored.
4. While $Q \neq \emptyset$, repeat the following:
 - (4.1) Pop a vertex v from Q .
 - (4.2) For each unexplored neighbor w of v , with probability \hat{p} add w to Q and mark w as colored. Mark w as explored.
5. Let $S \in \mathcal{S}(G)$ be the graphlet corresponding to the colored connected component of v . Accept S with probability:

$$\frac{1}{|S|} \cdot (1 - \hat{p})^{(\Delta-2)|S|+2-|\partial S|} \left(\frac{\lambda}{\hat{\lambda}} \right)^{|S|}.$$

6. If S is rejected, go back to Step 2 and repeat.

The analysis of this algorithm is similar to that in the proof of Theorem 8. Let $n = |V|$. The probability that the algorithm outputs S in an iteration is:

$$\sum_{v \in S} \frac{1}{n} \cdot \hat{p}^{|S|} (1 - \hat{p})^{|\partial S|} \cdot \frac{1}{|S|} \cdot (1 - \hat{p})^{(\Delta-2)|S|+2-|\partial S|} \left(\frac{\lambda}{\hat{\lambda}} \right)^{|S|} = \frac{(1 - \hat{p})^2 \lambda^{|S|}}{n}. \quad (16)$$

Hence, conditioned on acceptance, the probability of obtaining $S \in \mathcal{S}(G)$ is thus $\nu_{G,\lambda}(S)$, and so the output distribution of the algorithm is $\nu_{G,\lambda}$.

For the running time of the algorithm, we note that Step 4 of the algorithm is analogous to Step 3 of the algorithm in the proof of Theorem 8, and so the expected running time of each round is also bounded by a constant $C = C(\Delta, \hat{p}) > 0$. Let T be the overall the running time of the algorithm. From (16), we have that the overall acceptance probability in a round is $\rho = \frac{(1-\hat{p})^2 Z(G, \lambda)}{n}$. Then, as in (15), we deduce that $\mathbb{E}[T] = O(nZ(G, \lambda)^{-1})$. Since $Z(G, \lambda) \geq n\lambda$, we have $\mathbb{E}[T] = O(1)$. ◀

3 Applications to Polymer Models

In this section, we show how to use our algorithm for sampling rooted vertex-labeled graphlets from Section 2 to sample from subset polymer models and prove Theorem 4.

Consider a subset polymer model on an n -vertex graph $G = (V, E)$; see Section 1.2 for the definition. Recall that we use \mathcal{C}_v for the set of all polymers containing vertex $v \in V$, and let γ_\emptyset denote the empty polymer. Define the distribution ν_v on $\mathcal{C}_v \cup \{\gamma_\emptyset\}$ by

$$\nu_v(\gamma) = \frac{w_\gamma}{\sum_{\hat{\gamma} \in \mathcal{C}_v \cup \{\gamma_\emptyset\}} w_{\hat{\gamma}}},$$

where we assign $w_{\gamma_\emptyset} = 1$. The following Markov chain on $\Omega(\mathcal{C})$, introduced in [16], has stationary distribution μ and mixes rapidly in $O(n \log n)$ steps under the polymer mixing condition (6).

Polymer dynamics. Given a configuration $X_t \in \Omega(\mathcal{C})$, form X_{t+1} as follows:

1. Pick $v \in V$ uniformly at random and let $S_v = \{\gamma \in X_t : v \in \gamma\}$ (note that S_v is either empty or contains 1 polymer).
2. With probability 1/2, let $X_{t+1} = X_t \setminus S_v$.
3. With probability 1/2 (exclusively of step 2), sample γ from ν_v . Let $X_{t+1} = X_t \cup \{\gamma\}$ if $X_t \cup \{\gamma\} \in \Omega(\mathcal{C})$ and let $X_{t+1} = X_t$ otherwise.

To implement a single update step, one must sample from ν_v in Step 3. To do so efficiently, in [16] the much stricter polymer sampling condition (7) was required. Here we give a fast perfect sampler for ν_v under a much weaker condition.

► **Theorem 10.** *Consider a subset polymer model on a family of n -vertex graphs of maximum degree $\Delta \geq 3$ with computationally feasible weights that satisfy $w_\gamma \leq \lambda^{|\gamma|}$ for some $\lambda < \lambda_*(\Delta, q)$. There is a randomized algorithm to sample perfectly from ν_v for any $v \in V$ with expected running time bounded by a function of λ, Δ , and q .*

Proof. This follows from Theorem 8. ◀

Using this theorem and the fast mixing result for the polymer dynamics of [16], one can approximately sample from μ whenever both the polymer mixing condition (6) and the assumptions of Theorem 10 hold. We further improve this by giving a perfect sampling algorithm that works whenever a new condition (4) is satisfied (our algorithm also requires the assumptions in Theorem 10). In all known examples, condition (4) is more permissive than the polymer mixing condition (6) from [16].

3.1 Perfect Sampling for polymer systems: Proof of Theorem 4

As mentioned, the polymer dynamics from [16] is not a perfect sampler. We propose here a different algorithm to output a perfect sample from μ . Our algorithm uses the coupling from the past method [44] and the notion of bounding Markov chains [33, 30] to efficiently implement it.

We proceed with the proof of Theorem 4. We start with the description of a grand coupling for the polymer dynamics, which is then used to implement a coupling from the past algorithm. For an n -vertex graph $G = (V, E)$, let $\{X_t^\Gamma\}$ denote an instance of the polymer dynamics started from the polymer configuration $\Gamma \in \Omega(\mathcal{C})$. For all $\Gamma \in \Omega(\mathcal{C})$, the chains $\{X_t^\Gamma\}$ are coupled by choosing the same uniform random vertex $v \in V$, the same polymer γ sampled from ν_v , and the same uniform random number in $[0, 1]$ to decide whether to remove S_v (Step 2) or to add γ (Step 3). A coupling from the past algorithm will find a time $-T$ such the grand coupling started from all possible states at time $-T$ coalesces to a single state by time 0. This guarantees that the output of the algorithm, that is the state at time 0, has distribution μ (see Theorem 1 from [44]). Such a T can be found with a binary search procedure. Unfortunately, implementing the coupling from the past algorithm in this manner for the polymer dynamics Markov chain is infeasible in our setting, since it requires simulating an exponential number of copies of the polymer dynamics, one from each $\Gamma \in \Omega(\mathcal{C})$.

To work around this, we consider a bounding Markov chain for the polymer dynamics rather than the polymer dynamics chain itself. Bounding Markov chains were pioneered in [33, 30] as a method for efficiently implementing coupling from the past. The bounding chain for the polymer dynamics has state space $\Omega(\mathcal{C}) \times 2^{\mathcal{C}}$ and will be denoted by $\{B_t, D_t\}$, where $B_t \in \Omega(\mathcal{C})$ and $D_t \subseteq \mathcal{C}$ are sets of polymers. The chain will maintain throughout that all polymers in B_t are compatible and that every polymer in B_t is compatible with every polymer in D_t . The polymers in D_t do *not* need to be compatible with each other. A step of the bounding Markov chain is defined next.

Polymer Dynamics Bounding Chain. Given $\{B_t, D_t\}$, the chain generates $\{B_{t+1}, D_{t+1}\}$ by:

1. Uniformly at random, select $v \in V$.
2. With probability $1/2$, remove all polymers containing v by setting $B_{t+1} = B_t \setminus \mathcal{C}_v$ and $D_{t+1} = D_t \setminus \mathcal{C}_v$.
3. With the remaining probability $1/2$, draw a sample γ according to ν_v and:
 - a. If γ is compatible with B_t and γ is compatible with $D_t \setminus \{\gamma\}$, let $B_{t+1} = B_t \cup \{\gamma\}$ and let $D_{t+1} = D_t \setminus \{\gamma\}$.
 - b. Else if γ is compatible with B_t but γ is not compatible with D_t , let $B_{t+1} = B_t$ and let $D_{t+1} = D_t \cup \{\gamma\}$.
 - c. If γ is incompatible with B_t , do nothing: $B_{t+1} = B_t$ and $D_{t+1} = D_t$.

Observe that polymers are only added to B_t if they are compatible with all other polymers in B_t ; hence, if B_0 is a valid polymer configuration, so is B_t for all $t \geq 0$.

To implement a step of the polymer dynamics bounding chain it suffices to pick a vertex $v \in V$ uniformly at random, a uniform random number in $[0, 1]$, and a polymer γ from ν_v , just like for the polymer dynamics. Hence, we can couple the evolution of $\{B_t, D_t\}$ with the grand coupling of the polymer dynamics described earlier. If we set $B_0 = \emptyset$ and $D_0 = \mathcal{C}$, it can be checked that for all $\Gamma \in \Omega(\mathcal{C})$ and all $t \geq 0$:

$$B_t \subseteq X_t^\Gamma \subseteq B_t \cup D_t.$$

Indeed, this holds initially for $t = 0$, and the grand coupling ensures that whenever a polymer is removed from X_t^Γ it is also removed from B_t , and whenever a polymer is added to X_t^Γ it is also added to B_t or D_t . Consequently, $\{B_t, D_t\}$ is a bounding chain for the polymer dynamics. In particular, the first time $B_t = B_t \cup D_t$, all instances X_t^Γ have necessarily coalesced to the same configuration. This bounding chain allows us to implement the coupling from the past algorithm efficiently, as follows.

Coupling from the Past. Set $k = 1$.

- (A) For $t = -2^k, -2^k + 1, \dots, -2^{k-1}$ generate $\rho_t = (v_t, \gamma_t, r_t)$ by choosing $v_t \in V$ uniformly at random, $r_t \in [0, 1]$ uniformly at random, and by sampling $\gamma_t \in \mathcal{C}_{v_t}$ from ν_{v_t} .
- (B) Set $B_{-2^k} = \emptyset$ and $D_{-2^k} = \mathcal{C}$.
- (C) Simulate the polymer dynamics bounding chain from time -2^k to time 0 using $\rho_{-2^k}, \dots, \rho_{-1}$.
- (D) If $B_0 = B_0 \cup D_0$, then output B_0 ; otherwise set $k \rightarrow k + 1$ and repeat the process from Step (A).

This implementation of the coupling from the past algorithm provides a perfect sample from μ ; see [44]. It remains for us to show that it can be efficiently implemented. For this, we show first that the expected number of steps of the polymer dynamics bounding chain throughout the execution of the algorithm is $O(n \log n)$. Afterwards, we will show how to implement steps so that they can be executed in amortized constant expected time.

► **Lemma 11.** *Suppose the subset polymer model satisfies condition (4). Then, the expected number of steps of the polymer dynamics bounding chain in the coupling from past algorithm is $O(n \log n)$.*

Proof. See the full version [6], where a potential ϕ_t describing the size of D_t is introduced and shown to decrease quickly. ◀

It remains to consider how to efficiently implement the steps of the polymer dynamics bounding chain. This is subtle because D_t may initially contain an exponentially large number of polymers, and care is thus needed in how D_t is represented and stored. The following lemma says we can represent and update B_t and D_t efficiently.

► **Lemma 12.** *There exists a compact representation of B_t and D_t that uses $O(n + t)$ space in expectation. Using this representation, each iteration of the Polymer Dynamics Bounding Chain can be executed in amortized constant expected time.*

Proof. See the full version [6]. ◀

References

- 1 Matteo Agostini, Marco Bressan, and Shahrzad Haddadan. Mixing time bounds for graphlet random walks. *Information Processing Letters*, 152:105851, 2019.
- 2 Konrad Anand and Mark Jerrum. Perfect sampling in infinite spin systems via strong spatial mixing. *arXiv preprint*, 2021. [arXiv:2106.15992](https://arxiv.org/abs/2106.15992).
- 3 Nima Anari, Kuikui Liu, and Shayan Oveis Gharan. Spectral independence in high-dimensional expanders and applications to the hardcore model. *SIAM Journal on Computing: Special Section, FOCS 2020*, pages FOCS20:1–FOCS20:37, 2022.
- 4 Kim Baskerville, Peter Grassberger, and Maya Paczuski. Graph animals, subgraph sampling, and motif search in large networks. *Physical Review E*, 76(3):036107, 2007.
- 5 Mansurul A Bhuiyan, Mahmudur Rahman, Mahmuda Rahman, and Mohammad Al Hasan. Guise: Uniform sampling of graphlets for large graph analysis. In *2012 IEEE 12th International Conference on Data Mining*, pages 91–100. IEEE, 2012.
- 6 Antonio Blanca, Sarah Cannon, and Will Perkins. Fast and perfect sampling of subgraphs and polymer systems. *arXiv preprint*, 2022. [arXiv:2202.05907](https://arxiv.org/abs/2202.05907).
- 7 Christian Borgs, Jennifer Chayes, Tyler Helmuth, Will Perkins, and Prasad Tetali. Efficient sampling and counting algorithms for the Potts model on \mathbb{Z}^d at all temperatures. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing (STOC)*, pages 738–751, 2020.

- 8 Christian Borgs, Jennifer Chayes, Jeff Kahn, and László Lovász. Left and right convergence of graphs with bounded degree. *Random Structures & Algorithms*, 42(1):1–28, 2013.
- 9 Marco Bressan. Efficient and near-optimal algorithms for sampling connected subgraphs. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing (STOC)*, pages 1132–1143, 2021.
- 10 Marco Bressan, Flavio Chierichetti, Ravi Kumar, Stefano Leucci, and Alessandro Panconesi. Counting graphlets: Space vs time. In *Proceedings of the tenth ACM International Conference on Web Search and Data Mining (WSDM)*, pages 557–566, 2017.
- 11 Marco Bressan, Flavio Chierichetti, Ravi Kumar, Stefano Leucci, and Alessandro Panconesi. Motif counting beyond five nodes. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 12(4):1–25, 2018.
- 12 Marco Bressan, Stefano Leucci, and Alessandro Panconesi. Faster motif counting via succinct color coding and adaptive sampling. *ACM Transactions on Knowledge Discovery from Data (TKDD)*, 15(6):1–27, 2021.
- 13 Sarah Cannon and Will Perkins. Counting independent sets in unbalanced bipartite graphs. In *Proceedings of the Thirty-First Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1456–1466, 2020.
- 14 Charlie Carlson, Ewan Davies, and Alexandra Kolla. Efficient algorithms for the Potts model on small-set expanders. *arXiv preprint*, 2020. [arXiv:2003.01154](https://arxiv.org/abs/2003.01154).
- 15 Xiaowei Chen, Yongkun Li, Pinghui Wang, and John CS Lui. A general framework for estimating graphlet statistics via random walk. *Proceedings of the VLDB Endowment*, 10(3), 2016.
- 16 Zongchen Chen, Andreas Galanis, Leslie A Goldberg, Will Perkins, James Stewart, and Eric Vigoda. Fast algorithms at low temperatures via Markov chains. *Random Structures & Algorithms*, 58(2):294–321, 2021.
- 17 Zongchen Chen, Andreas Galanis, Daniel Štefankovič, and Eric Vigoda. Sampling colorings and independent sets of random regular bipartite graphs in the non-uniqueness region. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2198–2207. SIAM, 2022.
- 18 Zongchen Chen, Kuikui Liu, and Eric Vigoda. Rapid mixing of Glauber dynamics up to uniqueness via contraction. In *2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS)*, pages 1307–1318. IEEE, 2020.
- 19 Zongchen Chen, Kuikui Liu, and Eric Vigoda. Optimal mixing of Glauber dynamics: Entropy factorization via high-dimensional expansion. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing (STOC)*, pages 1537–1550, 2021.
- 20 Matthew Coulson, Ewan Davies, Alexandra Kolla, Viresh Patel, and Guus Regts. Statistical physics approaches to unique games. In *35th Computational Complexity Conference (CCC)*, 2020.
- 21 Martin Dyer, Leslie Ann Goldberg, Catherine Greenhill, and Mark Jerrum. The relative complexity of approximate counting problems. *Algorithmica*, 38(3):471–500, 2004.
- 22 Sacha Friedli and Yvan Velenik. *Statistical mechanics of lattice systems: a concrete mathematical introduction*. Cambridge University Press, 2017.
- 23 Tobias Friedrich, Andreas Göbel, Martin S Krejca, and Marcus Pappik. Polymer dynamics via cliques: New conditions for approximations. *arXiv preprint*, 2020. [arXiv:2007.08293](https://arxiv.org/abs/2007.08293).
- 24 Andreas Galanis, Leslie Ann Goldberg, and James Stewart. Fast algorithms for general spin systems on bipartite expanders. *ACM Transactions on Computation Theory (TOCT)*, 13(4):1–18, 2021.
- 25 Andreas Galanis, Leslie Ann Goldberg, and James Stewart. Fast mixing via polymers for random graphs with unbounded degree. *Information and Computation*, 285:104894, 2022.
- 26 Andreas Galanis, Daniel Štefankovič, and Eric Vigoda. Inapproximability of the partition function for the antiferromagnetic Ising and hard-core models. *Combinatorics, Probability and Computing*, 25(4):500–559, 2016.

- 27 Oded Goldreich and Dana Ron. Property testing in bounded degree graphs. In *Proceedings of the twenty-ninth annual ACM symposium on Theory of computing (STOC)*, pages 406–415, 1997.
- 28 Joshua A Grochow and Manolis Kellis. Network motif discovery using subgraph enumeration and symmetry-breaking. In *Annual International Conference on Research in Computational Molecular Biology*, pages 92–106. Springer, 2007.
- 29 Christian Gruber and Hervé Kunz. General properties of polymer systems. *Communications in Mathematical Physics*, 22(2):133–161, 1971.
- 30 Olle Haggstrom and Karin Nelander. On exact simulation of Markov random fields using coupling from the past. *Scandinavian Journal of Statistics*, 26(3):395–411, 1999.
- 31 Tyler Helmuth, Matthew Jenssen, and Will Perkins. Finite-size scaling, phase coexistence, and algorithms for the random cluster model on random graphs. *Annales de l’Institut Henri Poincaré B*. To appear, 2022.
- 32 Tyler Helmuth, Will Perkins, and Guus Regts. Algorithmic Pirogov–Sinai theory. *Probability Theory and Related Fields*, 176(3):851–895, 2020.
- 33 Mark Huber. Perfect sampling using bounding chains. *The Annals of Applied Probability*, 14(2):734–753, 2004.
- 34 Svante Janson, Andrzej Rucinski, and Tomasz Luczak. *Random graphs*. John Wiley & Sons, 2011.
- 35 Matthew Jenssen, Peter Keevash, and Will Perkins. Algorithms for #BIS-hard problems on expander graphs. *SIAM Journal on Computing*, 49(4):681–710, 2020.
- 36 Matthew Jenssen, Aditya Potukuchi, and Will Perkins. Approximately counting independent sets in bipartite graphs via graph containers. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 499–516. SIAM, 2022.
- 37 Madhav Jha, C Seshadhri, and Ali Pinar. Path sampling: A fast and provable method for estimating 4-vertex subgraph counts. In *Proceedings of the 24th International Conference on World Wide Web (WWW)*, pages 495–505, 2015.
- 38 Nadav Kashtan, Shalev Itzkovitz, Ron Milo, and Uri Alon. Efficient sampling algorithm for estimating subgraph concentrations and detecting network motifs. *Bioinformatics*, 20(11):1746–1758, 2004.
- 39 Roman Kotecký and David Preiss. Cluster expansion for abstract polymer models. *Communications in Mathematical Physics*, 103(3):491–498, 1986.
- 40 Chao Liao, Jiabao Lin, Pinyan Lu, and Zhenyu Mao. Counting independent sets and colorings on random regular bipartite graphs. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2019)*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2019.
- 41 Xuesong Lu and Stéphane Bressan. Sampling connected induced subgraphs uniformly at random. In *International Conference on Scientific and Statistical Database Management*, pages 195–212. Springer, 2012.
- 42 Ryuta Matsuno and Aristides Gionis. Improved mixing time for k -subgraph sampling. In *Proceedings of the 2020 SIAM International Conference on Data Mining*, pages 568–576. SIAM, 2020.
- 43 Kirill Paramonov, Dmitry Shemetov, and James Sharpnack. Estimating graphlet statistics via lifting. In *Proceedings of the 25th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining*, pages 587–595, 2019.
- 44 James Gary Propp and David Bruce Wilson. Exact sampling with coupled Markov chains and applications to statistical mechanics. *Random Structures & Algorithms*, 9(1-2):223–252, 1996.
- 45 Andrew Read-McFarland and Daniel Štefankovič. The hardness of sampling connected subgraphs. In *Latin American Symposium on Theoretical Informatics*, pages 464–475. Springer, 2021.

4:18 Fast and Perfect Sampling of Subgraphs and Polymer Systems

- 46 Nino Shervashidze, SVN Vishwanathan, Tobias Petri, Kurt Mehlhorn, and Karsten Borgwardt. Efficient graphlet kernels for large graph comparison. In *Artificial intelligence and statistics*, pages 488–495. PMLR, 2009.
- 47 Allan Sly. Computational transition at the uniqueness threshold. In *2010 IEEE 51st Annual Symposium on Foundations of Computer Science (FOCS)*, pages 287–296. IEEE, 2010.
- 48 Allan Sly and Nike Sun. The computational hardness of counting in two-spin models on d -regular graphs. In *2012 IEEE 53rd Annual Symposium on Foundations of Computer Science*, pages 361–369. IEEE, 2012.
- 49 Dror Weitz. Counting independent sets up to the tree threshold. In *Proceedings of the thirty-eighth annual ACM symposium on Theory of computing*, pages 140–149, 2006.