

Sampling repulsive Gibbs point processes using random graphs

Tobias Friedrich* Andreas Göbel* Maximilian Katzmann†
 Martin S. Krejca‡ Marcus Pappik*

*Hasso Plattner Institute, University of Potsdam, Potsdam, Germany
 {tobias.friedrich, andreas.goebel, marcus.pappik}@hpi.de

† Karlsruhe Institute of Technology, Karlsruhe, Germany
 maximilian.katzmann@kit.edu

‡ Laboratoire d'Informatique (LIX), CNRS, École Polytechnique, Institut Polytechnique de Paris
 martin.krejca@polytechnique.edu

Abstract

We study computational aspects of repulsive Gibbs point processes, which are probabilistic models of interacting particles in a finite-volume region of space. We introduce an approach for reducing a Gibbs point process to the hard-core model, a well-studied discrete spin system. Given an instance of such a point process, our reduction generates a random graph drawn from a natural geometric model. We show that the partition function of a hard-core model on graphs generated by the geometric model concentrates around the partition function of the Gibbs point process. Our reduction allows us to use a broad range of algorithms developed for the hard-core model to sample from the Gibbs point process and approximate its partition function. This is, to the extend of our knowledge, the first approach that deals with pair potentials of unbounded range. We compare the resulting algorithms with recently established results and study further properties of the random geometric graphs with respect to the hard-core model.

1 Introduction

Gibbs point processes are a tool for modelling a variety of phenomena that can be described as distributions of random spatial events [6, 36]. Such phenomena include the location of stars in the universe, a sample of cells under the microscope, or the location of pores and cracks in the ground (see [37, 45] for more on applications of Gibbs point processes). In statistical physics, such point processes are frequently used as stochastic models for gases or liquids of interacting particles [39].

A Gibbs point process on a finite-volume region V is parameterized by a fugacity λ and a pair potential ϕ expressing the interactions between pairs of points. Every point configuration in the region is assigned a weight according to the pair interactions ϕ between all pairs of points in the configuration. One can then think of a Gibbs point process as a Poisson point process of intensity λ , where the density of each configuration is scaled proportionally to its weight. The density is normalized by the partition function, which is the integral of the weights over the configuration space (see [Section 1.1](#) for a formal definition of the model). The most famous example of such a process is the *hard-sphere model*, a model of a random

packing of equal-sized spheres with radius r . The pair potential in the hard-sphere model defines hard-core interactions, i.e., configurations where two points closer than some distance $2r$ have weight zero, while all other configurations have weight one. In this article, we consider Gibbs point processes with *repulsive* potentials, that is, pair potentials in which adding a point to a configuration does not increase its weight. The hard-sphere model, for example, does have a repulsive pair potential, however, we do not restrict ourselves to hard-core potentials and allow for soft-core interactions.

The two most fundamental algorithmic tasks considered on Gibbs point processes are to sample from the model and to compute its partition function, which are closely related. Understanding for which potentials and fugacities these two tasks are tractable is an ambitious endeavour. Towards this goal, there has been a plethora of algorithmic results on Gibbs point processes spanning several decades. Notably, the Markov chain Monte Carlo method was developed for sampling an instance of the hard-sphere model with 224 particles [32]. Since then, a variety of exact and approximate sampling algorithms for such point processes have been proposed in the literature, and their efficiency has been studied extensively both without [19, 22] and with rigorous running time guarantees [35, 24, 21, 34, 1]. The key objective of rigorous works is to identify a parameter regime for their respective model for which a randomized algorithm for sampling and approximating the partition function exists with running time polynomial in the volume of the region. In addition, deterministic algorithms for approximating the partition function have also appeared in the literature [17, 26] with running time quasi-polynomial in the volume of the region.

This recent flurry of algorithmic results on Gibbs point processes can be attributed to progress in understanding the computational properties of discrete spin systems, such as the hard-core model. Within these works, two main approaches can be identified for transferring insights from discrete spin systems to Gibbs point processes. The first one, which includes results such as [23, 34, 1], considers properties proven to hold in discrete spin systems and translates them to the continuous setting of Gibbs point processes. More precisely, these works consider the notion of strong spatial mixing, which has been strongly connected to algorithmic properties of discrete spin systems [47, 40, 16], and translate it to an analogous notion for Gibbs point processes to obtain algorithmic results. A common pattern in these works is that once the parameter regime for which strong spatial mixing holds is established, one needs to prove from scratch that this implies efficient algorithms in Gibbs point processes. In addition, the definition of strong spatial mixing for Gibbs points processes assumes that the pair interactions of two particles is always of bounded range, i.e., if two particles are placed at distance greater than some constant $r \in \mathbb{R}_{\geq 0}$, they do not interact with each other.

The second approach, used in [18, 17], is to discretize the model, i.e., reduce it to an instance of the hard-core model and then solve the respective algorithmic problem for the hard-core model. In this case, the algorithmic arsenal developed over the years for the hard-core model is now readily available for the instances resulting from this reduction. The main downside of these approaches is that they only apply to the hard-sphere model, a special case of bounded-range repulsive interactions.

Our contributions. We introduce a natural approach for reducing repulsive point processes to the hard-core model. Given an instance $(\mathbb{V}, \lambda, \phi)$ of such a point process, we generate a random graph by sampling $n \in \Theta(\mathbb{V}^2)$ point-vertices independently and uniformly at random in \mathbb{V} and by connecting each pair of points with an edge drawn with an appropriate probability, which depends on ϕ . We show that computational properties of the hard-core model on graphs generated by this model and with an appropriately scaled fugacity transfer to the originating Gibbs point process.

We first show that the partition function of the hard-core model on these graphs concentrates around the partition function of the Gibbs point process. Using existing algorithms for the hard-core model as a

black box, our result immediately yields randomized approximation algorithms for the partition function of the point process in running time polynomial in the volume of \mathbb{V} . Furthermore, we show that sampling an independent set from the generated hard-core model and returning the positions of its vertices in \mathbb{V} results in an approximate sampler from the distribution of the Gibbs point process. Our approach, in contrast to all previous algorithmic work in the literature, does not require the pair potential ϕ of the point process to be of bounded range. This includes various models of interest in statistical physics, such as the (hard-core) Yukawa model [15, 38], the Gaussian overlap model [7], the generalized exponential model [5], and the Yoshida–Kamakura model [48].

Finally, we identify the parameter regime (in terms of λ and ϕ) for which the generated hard-core instance exhibits strong spatial mixing. This parameter regime is identical to the best known regime for which repulsive point processes with bounded-range potentials have strong spatial mixing [34].

1.1 Gibbs point processes

We now formally define the notion of a Gibbs point process. As usual in the theory of point processes, we assume the underlying space to be a complete, separable metric space (\mathbb{X}, d) equipped with the Borel algebra $\mathcal{B} = \mathcal{B}(\mathbb{X})$ and a reference measure ν on $(\mathbb{X}, \mathcal{B})$ that assigns finite volume to bounded measurable sets. In this work, we study Gibbs point processes $P_{\mathbb{V}}^{(\lambda, \phi)}$ on bounded measurable regions $\mathbb{V} \subseteq \mathbb{X}$ that are parameterized by a *fugacity* parameter $\lambda \in \mathbb{R}_{\geq 0}$ and a repulsive (i.e., non-negative), symmetric, measurable *potential function* $\phi: \mathbb{X}^2 \rightarrow \mathbb{R}_{\geq 0} \cup \{\infty\}$. Such a process $P_{\mathbb{V}}^{(\lambda, \phi)}$ is defined by a density with respect to a Poisson point process Q_{λ} with intensity λ on \mathbb{X} . For every finite point configuration $\mathbf{x} = (x_1, \dots, x_k) \in \mathbb{V}^k$, this density is proportional to $e^{-H(x_1, \dots, x_k)}$, where H is the *Hamiltonian*

$$H(x_1, \dots, x_k) = \sum_{\{i, j\} \in \binom{[k]}{2}} \phi(x_i, x_j).$$

More precisely, the density can be expressed explicitly as

$$\frac{dP_{\mathbb{V}}^{(\lambda, \phi)}}{dQ_{\lambda}}(x_1, \dots, x_k) = \frac{\mathbb{1}_{\forall i \in [k]: x_i \in \mathbb{V}} \cdot e^{-H(x_1, \dots, x_k)} e^{\lambda \nu(\mathbb{V})}}{\Xi_{\mathbb{V}}(\lambda, \phi)},$$

where the normalizing constant $\Xi_{\mathbb{V}}(\lambda, \phi)$ is the *partition function*

$$\Xi_{\mathbb{V}}(\lambda, \phi) = 1 + \sum_{k \in \mathbb{N}_{\geq 1}} \frac{\lambda^k}{k!} \int_{\mathbb{V}^k} e^{-H(x_1, \dots, x_k)} \nu^k(d\mathbf{x}).$$

1.2 Randomized reduction to the hard-core model

Our approach is to reduce the problem of sampling from a repulsive Gibbs point process and approximating its partition function to the analogous problems for a discrete hard-core model, which we briefly introduce. For an undirected graph $G = (V, E)$, let $\mathcal{I}(G) \subseteq 2^V$ denote the set of independent sets of G . For a vertex activity $\gamma \in \mathbb{R}_{\geq 0}$, the *hard-core model* on G is a probability distribution $\mu_G^{(\gamma)}$ on $\mathcal{I}(G)$ that assigns each independent set $I \in \mathcal{I}(G)$ a probability proportional to $\gamma^{|I|}$. The normalizing constant of this distribution, $Z_G(\gamma) = \sum_{I \in \mathcal{I}(G)} \gamma^{|I|}$, is called the *hard-core partition function* on G .

The goal is to reduce the problem of approximate sampling from $P_{\mathbb{V}}^{(\lambda, \phi)}$ to approximate sampling from $\mu_G^{(\gamma)}$ and, similarly, to reduce the problem of approximating $\Xi_{\mathbb{V}}(\lambda, \phi)$ to approximating $Z_G(\gamma)$ for a suitably chosen graph G and vertex activity γ . The advantage of this approach is that sampling from a hard-core

model as well as approximating hard-core partition functions are well studied problems. Specifically, a sequence of recent papers [3, 4, 10, 11, 12] established approximate sampling from $\mu_G^{(\gamma)}$ in $\tilde{O}(|V|)$ running time and randomized approximation of $Z_G(\gamma)$ in $\tilde{O}(|V|^2)$ running time for graphs G with maximum degree d for all γ strictly below the *tree threshold* $\gamma_c(d) := \frac{(d-1)^{d-1}}{(d-2)^d}$.

Our reduction is inspired by the discretization schemes in [17, 18]. These approaches are limited to the hard-sphere model (and similar models with hard-core interactions) in specific regions of Euclidean space. In this setting, the utilized graph G is essentially a unit-disk graph in \mathbb{V} . This procedure comes with two major disadvantages. Firstly, the analysis heavily depends on geometric arguments and is therefore restricted to regions in Euclidean space that satisfy various requirements. Secondly, it is not obvious how this technique extends to general repulsive potentials ϕ and especially how to account for soft-core interactions.

We circumvent the above problems by investigating hard-core models on a suitably chosen family of random graphs. For a bounded measurable region $\mathbb{V} \subseteq \mathbb{X}$, let $u_{\mathbb{V}}$ denote the uniform distribution on \mathbb{V} . That is, for $x \sim u_{\mathbb{V}}$, we have $\Pr[x \in A] = \frac{\nu(A)}{\nu(\mathbb{V})}$ for every measurable $A \subseteq \mathbb{V}$, and $\Pr[x \notin \mathbb{V}] = 0$. For a repulsive potential ϕ and a positive integer $n \in \mathbb{N}_{\geq 1}$, we consider a random-graph model $\zeta_{\mathbb{V}, \phi}^{(n)}$ on the set of undirected graphs with vertex set $[n]$, where $\zeta_{\mathbb{V}, \phi}^{(n)}$ is defined by the following natural procedure to generate a graph:

1. For each $i \in [n]$, draw a uniform random point $x_i \sim u_{\mathbb{V}}$ independently.
2. For all $i, j \in [n]$ with $i \neq j$, connect i and j with an edge with probability $1 - e^{-\phi(x_i, x_j)}$ independently.

Readers familiar with graphons might notice that this random-graph model can be expressed as a graphon-based random graph (W -random graph) for a suitably chosen graphon W . We discuss this perspective later in the introduction. Moreover, we would like to mention that a similar graph construction based on points from a Poisson point process was used in [8] to prove uniqueness of the infinite-volume Gibbs measure for $\lambda < \frac{1}{C_\phi}$ via percolation.

The key property of the graphs from $\zeta_{\mathbb{V}, \phi}^{(n)}$ is that, for a suitably chosen vertex activity γ , their hard-core partition functions concentrate around $\Xi_{\mathbb{V}}(\lambda, \phi)$. This property is at the core of our reduction.

► **Theorem 1.1.** Let (\mathbb{X}, d) be a complete, separable metric space, let $\mathcal{B} = \mathcal{B}(\mathbb{X})$ be the Borel algebra, and let ν be a locally finite reference measure on $(\mathbb{X}, \mathcal{B})$. Let $\mathbb{V} \subseteq \mathbb{X}$ be bounded and measurable, let $\lambda \in \mathbb{R}_{\geq 0}$, and let $\phi: \mathbb{X}^2 \rightarrow \mathbb{R}_{\geq 0} \cup \{\infty\}$ be a symmetric repulsive potential. For all $\varepsilon \in (0, 1]$, $\delta \in (0, 1]$ and $n \geq 4\varepsilon^{-2}\delta^{-1} \max\{e^6 \lambda^2 \nu(\mathbb{V})^2, \ln(4\varepsilon^{-1})^2\}$, it holds that, for $G \sim \zeta_{\mathbb{V}, \phi}^{(n)}$,

$$\Pr\left[\left|Z_G\left(\frac{\lambda\nu(\mathbb{V})}{n}\right) - \Xi_{\mathbb{V}}(\lambda, \phi)\right| \geq \varepsilon\Xi_{\mathbb{V}}(\lambda, \phi)\right] \leq \delta. \quad \blacktriangleleft$$

Informally, [Theorem 1.1](#) says that, for $n \in \Theta(\nu(\mathbb{V})^2)$ and $G \sim \zeta_{\mathbb{V}, \phi}^{(n)}$, the hard-core partition function $Z_G(\gamma(n))$ with $\gamma(n) = \frac{\lambda\nu(\mathbb{V})}{n}$ is strongly concentrated around the partition function of the repulsive Gibbs point process $\Xi_{\mathbb{V}}(\lambda, \phi)$. In [17, Proposition 5.8], it was argued that the partition function of an unrestricted Poisson point process in a bounded measurable region \mathbb{V} of Euclidean space cannot be approximated by the hard-core partition function $Z_G(\gamma(n))$ for any graph G on n vertices if $n \in o(\nu(\mathbb{V})^2)$. As the unrestricted Poisson point process is a special case of a repulsive Gibbs point process with constant zero potential, this implies that our concentration result in [Theorem 1.1](#) is tight in terms of its asymptotic dependency on the volume $\nu(\mathbb{V})$.

A remarkable aspect of [Theorem 1.1](#) is the generality in which it holds with respect to the underlying space \mathbb{X} . In particular, we do not need any additional assumptions regarding its geometry. We achieve this by deriving [Theorem 1.1](#) from a corollary of the Efron–Stein inequality [14]. This corollary gives a convenient-to-use way for proving concentration of functions of independent random inputs, given that changing one input of the function only leads to small relative changes of its output. Using this corollary, we prove [Theorem 1.1](#) by bounding changes in the hard-core partition function of a graph that are caused by small alternations of the graph structure. We proceed by discussing this approach in detail.

Proving concentration

We prove [Theorem 1.1](#) in two steps. First, we show that, for $G \sim \zeta_{\mathbb{V},\phi}^{(n)}$ and $\gamma(n) = \frac{\lambda v(\mathbb{V})}{n}$, the expected hard-core partition function $\mathbb{E}[Z_G(\gamma(n))]$ converges rapidly to the partition function of the point process $\Xi_{\mathbb{V}}(\lambda, \phi)$ as n grows. We prove this by fixing $n \in \Theta(v(\mathbb{V})^2)$ sufficiently large and rewriting $\mathbb{E}[Z_G(\gamma(n))]$ as a sum of expectations. We then relate each of the first $m < n$ terms in that sum with the corresponding term in $\Xi_{\mathbb{V}}(\lambda, \phi)$, which requires comparing the expected number of independent sets of size $k \leq m$ in a graph $G \sim \zeta_{\mathbb{V},\phi}^{(n)}$ with $\frac{n^k}{k!v(\mathbb{V})^k} \int_{\mathbb{V}^k} e^{-H(\mathbf{x})} d\mathbf{x}$. Finally, we argue that all terms of order higher than m can be discarded.

Once convergence of the expectation is established, it remains to prove that the distribution of the partition functions $Z_G(\gamma(n))$ for $G \sim \zeta_{\mathbb{V},\phi}^{(n)}$ concentrates around this expectation. To prove the latter, we derive [Corollary 3.3](#) from the Efron–Stein inequality [14]. This corollary roughly states that the output of a function f on a product of probability spaces concentrates around its expectation if f exhibits sufficiently small *relative* changes when any component of its input is changed. Similar methods for proving concentration usually require the output of the function f to exhibit small *absolute* changes (i.e., f to be Lipschitz, see [30, 31]), which does not hold in our setting.

To apply this concentration bound, we need to express the partition function of a graph drawn from $\zeta_{\mathbb{V},\phi}^{(n)}$ as a function of independent random inputs. To this end, we model a random graph $G \sim \zeta_{\mathbb{V},\phi}^{(n)}$ based on n points $\mathbf{x} = (x_i)_{i \in [n]}$, each independently drawn from $u_{\mathbb{V}}$, and $\frac{n(n-1)}{2}$ independent random variables $\mathbf{y} = (y_{i,j})_{1 \leq i < j \leq n}$, each uniformly distributed on the real interval $[0, 1]$. Given the random vectors \mathbf{x} and \mathbf{y} , we construct a graph by connecting vertices $i < j$ by an edge if and only if $y_{i,j} \leq 1 - e^{-\phi(x_i, x_j)}$. Note that the resulting graph is distributed according to $\zeta_{\mathbb{V},\phi}^{(n)}$. Thus, we express the hard-core partition function on the random-graph model $\zeta_{\mathbb{V},\phi}^{(n)}$ as a function $f(\mathbf{x}, \mathbf{y})$ for \mathbf{x} and \mathbf{y} as described above. The effect of changing a component of \mathbf{y} is bounded by the relative change of the hard-core partition function when adding or removing an edge. On the other hand, the effect of changing a component of \mathbf{x} , say x_i , is bounded by considering the change of the hard-core partition function when altering the neighborhood of a single vertex i . Bounding both effects and applying [Corollary 3.3](#) yields the desired concentration result ([Theorem 1.1](#)).

In fact, a similar argument as above applies to a broad class of antiferromagnetic spin systems on graphon-based random graphs that contains our application as a special case. We discuss this more general setting after demonstrating the sampling and approximation results for Gibbs point processes that can be obtained from [Theorem 1.1](#).

1.3 Algorithmic implications

We proceed by showcasing some algorithmic results for repulsive Gibbs point processes that follow from [Theorem 1.1](#). More specifically, we focus on ε -approximate sampling from the point process and obtaining

a randomized ε -approximation of the partition function. Formally, the problem of ε -approximate sampling from $P_{\mathbb{V}}^{(\lambda, \phi)}$ is defined as producing a random point configuration with a distribution that has a total variation distance of at most ε to $P_{\mathbb{V}}^{(\lambda, \phi)}$. Analogously, the problem of ε -approximating $\Xi_{\mathbb{V}}(\lambda, \phi)$ is defined as computing some value $x \in \mathbb{R}$ such that $(1 - \varepsilon)\Xi_{\mathbb{V}}(\lambda, \phi) \leq x \leq (1 + \varepsilon)\Xi_{\mathbb{V}}(\lambda, \phi)$. Moreover, an algorithm is called a *randomized ε -approximation* if it outputs an ε -approximation of $\Xi_{\mathbb{V}}(\lambda, \phi)$ with probability at least $\frac{2}{3}$. The choice of the constant $\frac{2}{3}$ is rather arbitrary here, as the error probability can be made smaller than every $\delta \in \mathbb{R}_{>0}$ by taking the median of $O(\log(\delta^{-1}))$ independent runs, as long as the error probability of each run is some constant smaller than $\frac{1}{2}$. Furthermore, we consider an ε -approximate sampler and a (randomized) ε -approximation algorithm as efficient if their running time is polynomial in the volume $\nu(\mathbb{V})$ and in ε^{-1} .

Recent rigorous results establish bounds on the fugacity regime of different models for which these algorithmic problems can be solved efficiently. Often, these bounds are stated in terms of the *temperedness constant* C_{ϕ} , which is defined as

$$C_{\phi} = \sup_{x_1 \in \mathbb{X}} \int_{\mathbb{X}} \left| 1 - e^{-\phi(x_1, x_2)} \right| \nu(dx_2).$$

This value can be seen as measure for the strength of interactions between points.

Given our concentration result ([Theorem 1.1](#)), a straightforward idea for approximating $\Xi_{\mathbb{V}}(\lambda, \phi)$ is to sample a graph $G \sim \zeta_{\mathbb{V}, \phi}^{(n)}$ and try to approximate its hard-core partition function. A refined version of this procedure leads to the following theorem.

► **Theorem 1.2.** Let (\mathbb{X}, d) be a complete, separable metric space, let $\mathcal{B} = \mathcal{B}(\mathbb{X})$ be the Borel algebra, and let ν be a locally finite reference measure on $(\mathbb{X}, \mathcal{B})$. Let $\mathbb{V} \subseteq \mathbb{X}$ be bounded and measurable, let $\lambda \in \mathbb{R}_{\geq 0}$, and let $\phi: \mathbb{X}^2 \rightarrow \mathbb{R}_{\geq 0} \cup \{\infty\}$ be a symmetric repulsive potential. Assume there is a sampler for $\zeta_{\mathbb{V}, \phi}^{(n)}$ with running time $t_{\mathbb{V}, \phi}(n)$.

If $\lambda < \frac{\varepsilon}{C_{\phi}}$, then, for all $\varepsilon \in (0, 1]$, there is a randomized ε -approximation algorithm for $\Xi_{\mathbb{V}}(\lambda, \phi)$ with running time in $\tilde{O}(\nu(\mathbb{V})^4 \varepsilon^{-6}) + t_{\mathbb{V}, \phi}(\tilde{O}(\nu(\mathbb{V})^2 \varepsilon^{-2}))$. ◀

With respect to sampling from $P_{\mathbb{V}}^{(\lambda, \phi)}$, it is less obvious how [Theorem 1.1](#) can be utilized. However, under mild assumptions, we obtain an approximate sampler, based on [Theorem 1.1](#), by the following procedure: Sample an independent set $I \in \mathcal{I}(G)$ (approximately) from $\mu_G^{(Y(n))}$ and output the point configuration $\{x_i\}_{i \in I}$. Given that $P_{\mathbb{V}}^{(\lambda, \phi)}$ is simple, which means that drawing a point configuration that contains the same point multiple times has probability zero, a refined version of the approach sketched above leads to the following result.

► **Theorem 1.3.** Let (\mathbb{X}, d) be a complete, separable metric space, let $\mathcal{B} = \mathcal{B}(\mathbb{X})$ be the Borel algebra, and let ν be a locally finite reference measure on $(\mathbb{X}, \mathcal{B})$. Let $\mathbb{V} \subseteq \mathbb{X}$ be bounded and measurable, let $\lambda \in \mathbb{R}_{\geq 0}$, and let $\phi: \mathbb{X}^2 \rightarrow \mathbb{R}_{\geq 0} \cup \{\infty\}$ be a symmetric repulsive potential. Assume we can sample from the uniform distribution $u_{\mathbb{V}}$ in time $t_{\mathbb{V}}$ and, for every $x, y \in \mathbb{V}$, we can evaluate $\phi(x, y)$ in time t_{ϕ} .

If the Gibbs point process $P_{\mathbb{V}}^{(\lambda, \phi)}$ is simple and $\lambda < \frac{\varepsilon}{C_{\phi}}$, then, for every $\varepsilon \in \mathbb{R}_{>0}$, there exists an ε -approximate sampling algorithm for $P_{\mathbb{V}}^{(\lambda, \phi)}$ with running time in $\tilde{O}(\nu(\mathbb{V})^2 \varepsilon^{-4} + \nu(\mathbb{V})^2 \varepsilon^{-3} t_{\mathbb{V}} + \nu(\mathbb{V})^4 \varepsilon^{-6} t_{\phi})$. ◀

There are two main differences in the assumptions of the approximation result ([Theorem 1.2](#)) and the sampling result ([Theorem 1.3](#)). First, the sampling result requires the point process to be simple. The

reason is that, in order to bound the total variation distance between the output of our sampler and $P_{\mathbb{V}}^{(\lambda, \phi)}$, we derive a density of that output with respect to a Poisson point process. This task is greatly simplified by assuming that $P_{\mathbb{V}}^{(\lambda, \phi)}$ is simple, as it allows for an easier characterization of the output distribution of our sampling, based on a theorem by Rényi–Mönch (see [13, Theorem 9.2.XII]). However, assuming the point process to be simple is only a minor restriction, as it is satisfied for most applications of point processes. For example, it is trivially satisfied if the reference volume measure ν is not-atomic (i.e., assigns volume 0 to single points). This includes the most frequently studied case of Gibbs point processes in Euclidean space but also a variety of other spaces, such as Gibbs point processes in hyperbolic spaces or in Riemannian manifolds.

Second, our sampling result requires efficient sampling from the uniform distribution $u_{\mathbb{V}}$ and an efficient way to compute the potential ϕ . In contrast to that, [Theorem 1.2](#) only assumes an efficient way to sample a graph from $\zeta_{\mathbb{V}, \phi}^{(n)}$. We state the theorems in this way to emphasize that, for approximating $\Xi_{\mathbb{V}}(\lambda, \phi)$, we only need to sample from the random-graph model $\zeta_{\mathbb{V}, \phi}^{(n)}$. Our sampling procedure additionally requires the position $x_i \in \mathbb{V}$ for each vertex $i \in [n]$ along with the graph to output the point configuration, associated to a random independent set drawn from the hard-core model.

Last, we briefly discuss the origin of the fugacity bound $\frac{e}{C_{\phi}}$ in our algorithmic results. Write $\gamma(n) = \frac{\lambda \nu(\mathbb{V})}{n}$ for every $n \in \mathbb{N}_{\geq 1}$. Note that our algorithms rely on either an efficient approximation of the hard-core partition function $Z_G(\gamma(n))$ or an efficient approximate sampler for an independent set from $\mu_G^{(\gamma(n))}$ for a random graph $G \sim \zeta_{\mathbb{V}, \phi}^{(n)}$. As discussed earlier, such computational results are known for general graphs of maximum degree d as long as the parameter γ is below the corresponding tree threshold $\gamma_c(d)$. Observe that $\gamma_c(d) \approx \frac{e}{d}$ for large d . Thus, roughly speaking, we can perform the necessary computational tasks as long as $\gamma(n) = \frac{\lambda \nu(\mathbb{V})}{n} < \frac{e}{d_G}$, where d_G is the maximum degree of the graph G that was drawn from $\zeta_{\mathbb{V}, \phi}^{(n)}$. Equivalently, this is $\lambda < \frac{en}{d_G \nu(\mathbb{V})}$. The main observation is now that, for $G \sim \zeta_{\mathbb{V}, \phi}^{(n)}$, the expected degree of an arbitrary vertex of G is upper-bounded by $\frac{n C_{\phi}}{\nu(\mathbb{V})}$. By proving that, with sufficiently high probability, the maximum degree d_G is not much larger than this value, we obtain the desired bound of $\frac{e}{C_{\phi}}$.

1.4 Strong spatial mixing, connective constants, and improved bounds

A property closely tied to the existence of efficient algorithms for the hard-core model is strong spatial mixing. Strong spatial mixing describes a particular way how dependencies between distant vertices in a graph G decay. The definition is easiest stated in terms of the occupation probability of a vertex $v \in V_G$ (i.e., the probability that v is in the independent set, drawn from a hard-core model on G) conditioned on certain vertices being occupied or unoccupied. Given two such conditions that differ at some vertex set $S \subset V_G$, strong spatial mixing requires that the resulting difference in the occupation probability of every vertex v is exponentially small in the graph distance between v and S^1 (see [Definition 6.1](#) and [Remark 6.2](#) for more details).

In a seminal paper, Weitz [47] proved that, for a graph G of maximum degree d_G and vertex $v \in V_G$, we can construct a tree with root v such that the occupation probability of v in the tree is the same as in G . This continues to hold when conditioning on the state of other vertices by translating the condition to the tree appropriately. It follows that if this tree exhibits strong spatial mixing with respect to the root, then this property also holds for G . Moreover, in this case, a recursive computation on this tree can be used

¹ Often it is more convenient to work with the occupation ratio, which is the occupation probability divided by the probability of the vertex to be unoccupied. However, the resulting strong spatial mixing definitions are equivalent.

to approximate the occupation probability of v , which results in a sampling and approximation algorithm with running time $|V_G|^{\mathcal{O}(\log(d_G))}$, for the hard-core model at vertex activity up to $\gamma_c(d_G)$ on G .

Subsequently, this result was improved by a more refined analysis [42, 40]. We elaborate further. The tree used by Weitz, which we refer to as the *Weitz tree*, is a truncated version of a self-avoiding walk tree² (see Section 6.3 for a formal definition). This truncation accounts for the effect of cycles on the hard-core distribution in the original graph. In [42, 40], it was shown that strong spatial mixing results can be derived by bounding the connective constant Δ , which describes the growth of the Weitz tree (see Definition 6.4). In particular, it was shown that strong spatial mixing applies up to a vertex activity of $\gamma < \gamma_c(\Delta)$, which improves bounds derived from the maximum degree. Consequentially, this implies sampling and approximation algorithm for the hard-core model for this parameter regime with running time $|V_G|^{\mathcal{O}(\log(\Delta))}$.

Inspired by these results on the hard-core model, Michelen and Perkins [33] recently introduced the potential-weighted connective constant Δ_ϕ for repulsive Gibbs point processes (see Section 6.1 for a formal definition). It can be seen as an alternative to C_ϕ that is more sensitive to the structure of the underlying space \mathbb{X} . In particular, for any non-trivial potential, Δ_ϕ is strictly smaller than the temperedness constant. Moreover, it was shown in [34] that repulsive Gibbs point processes with bounded-range potentials exhibit a notion of strong spatial mixing up to a fugacity of $\lambda < e/\Delta_\phi$ (see [34, Definition 1]). This result was used to derive a polynomial-time approximate sampling algorithm and a randomized approximation algorithm for the partition function of bounded-range repulsive Gibbs point processes in the same fugacity regime.

Given a repulsive Gibbs point process with fugacity $\lambda < e/\Delta_\phi$, we show in Section 6 that the hard-core model that is obtained from our reduction exhibits strong spatial mixing with high probability. In particular our result holds without bounded-range assumptions. We prove this by using the results in [42, 40]. Towards this, we establish a rigorous connection between the potential-weighted connective constant of a repulsive Gibbs point process and the connective constant of a graph from $\zeta_{\mathbb{V},\phi}^{(n)}$. More precisely, we show that, for any $\epsilon > 0$ and $n \geq \Theta(\nu(\mathbb{V}))$, the connective constant of a graph from $\zeta_{\mathbb{V},\phi}^{(n)}$ is bounded by $e^\epsilon \frac{n}{\nu(\mathbb{V})} \Delta_\phi$ with probability at least $1 - \frac{1}{n}$ (see Theorem 6.7 for the formal statement). To obtain this result, we make use of the fact that the construction of the Weitz tree leaves some degree of freedom. That is, for the same graph G and root vertex $v \in V_G$, different Weitz trees can be constructed, which differ in how the self-avoiding walk tree is truncated. In our setting, we carefully need to choose the truncation based on the underlying location of vertices in \mathbb{V} . In particular, it is important for us to define the connective constant in terms of the Weitz tree and not the full self-avoiding walk tree. The latter would only yield a bound of $\frac{n}{\nu(\mathbb{V})} C_\phi$, which would be no improvement over the maximum degree of $G \sim \zeta_{\mathbb{V},\phi}^{(n)}$.

Given the above graphical interpretation of Δ_ϕ , we immediately obtain that, for all $\lambda < e/\Delta_\phi$, a hard-core model with vertex activity $\gamma(n) = \frac{\nu(\mathbb{V})}{n} \lambda$ exhibits strong spatial mixing on $G \sim \zeta_{\mathbb{V},\phi}^{(n)}$ with probability at least $1 - \frac{1}{n}$ (see Corollary 6.10 for the formal statement). This result holds for any repulsive potential, without bounded-range assumption. However, it should be noted that the strong spatial mixing is with respect to the graph distance and not the distance metric of the underlying space \mathbb{X} .

This strong spatial mixing result for $G \sim \zeta_{\mathbb{V},\phi}^{(n)}$ gives our reduction further algorithmic consequences. Using the deterministic algorithm for approximating that partition function of a hard-core model proposed by Weitz [47] (see also [42, 40]) and Theorem 1.1, our strong spatial mixing result yields a randomized approximation for the partition function of repulsive Gibbs point processes with arbitrary range potentials for $\lambda < e/\Delta_\phi$ with quasi-polynomial running time $\nu(\mathbb{V})^{\mathcal{O}(\ln(\nu(\mathbb{V})))}$. A similar result can be derived in the setting of approximate sampling.

² In [47] the tree is actually not truncated but certain vertices in the tree are fixed to be always occupied or unoccupied. However, this is equivalent to truncating the tree.

1.5 A more general concentration result: antiferromagnetic partition functions on graphon-based random graphs

So far, we discussed how concentration of hard-core partition functions $Z_G\left(\frac{\lambda v(\mathbb{V})}{n}\right)$ for random graphs $G \sim \zeta_{\mathbb{V}, \phi}^{(n)}$, stated in [Theorem 1.1](#) is obtained from [Corollary 3.3](#). However, as hinted earlier, a more general concentration result can be obtained for a large class of antiferromagnetic two-state spin systems on graphon-based random graph models. As such spin systems have been studied extensively [[28](#), [41](#), [43](#)], we believe this result to be of independent interest. In what follows, we outline this more general concentration result and show how [Theorem 1.1](#) follows as a special case of it.

We start by introducing the class of spin systems to which it applies. For an undirected graph $G = (V, E)$ with vertices V and edges $E \subseteq \binom{V}{2}$, we denote by Σ_G the set of all functions $\sigma: V \rightarrow \{0, 1\}$. To simplify notation, we assume $V = [n]$ for some $n \in \mathbb{N}$. A *two-state spin system* with parameters $\gamma, \beta_0, \beta_1 \in \mathbb{R}_{\geq 0}$ on G is a probability distribution $\mu_G^{(\gamma, \beta_0, \beta_1)}$ on Σ_G with

$$\mu_G^{(\gamma, \beta_0, \beta_1)}(\sigma) = \frac{\gamma^{|\sigma|_1} \beta_0^{m_G^{(0)}(\sigma)} \beta_1^{m_G^{(1)}(\sigma)}}{Z_G(\gamma, \beta_0, \beta_1)},$$

where $|\sigma|_1 = |\sigma^{-1}(1)|$ counts the number of vertices that are assigned to 1, $m_G^{(a)}(\sigma) = \sum_{\{i, j\} \in E} \mathbb{1}_{\sigma(i)=\sigma(j)=a}$ counts the number of edges with both endpoints assigned to $a \in \{0, 1\}$, and the normalizing constant $Z_G(\gamma, \beta_0, \beta_1)$ is the partition function

$$Z_G(\gamma, \beta_0, \beta_1) = \sum_{\sigma \in \Sigma_G} \gamma^{|\sigma|_1} \beta_0^{m_G^{(0)}(\sigma)} \beta_1^{m_G^{(1)}(\sigma)}.$$

A two-state spin system is *antiferromagnetic* if $\beta_0 \beta_1 \leq 1$. Our concentration result applies to antiferromagnetic two-state spin systems with $\beta_0 = 1$. In this case, we omit β_0 completely, write $\beta = \beta_1 \in [0, 1]$, and denote the partition function by $Z_G(\gamma, \beta)$.

Our concentration result for partition functions $Z_G(\gamma, \beta)$ applies to all graphon-based random-graph models. Here, we refer to graphons in the most general sense, as defined in [[29](#), Chapter 13]. That is, for a probability space $\mathcal{X} = (X, \mathcal{A}, \xi)$, a *graphon* is a symmetric function $W: X^2 \rightarrow [0, 1]$ that is measurable with respect to the product algebra $\mathcal{A}^2 = \mathcal{A} \otimes \mathcal{A}$. Note that, even though we call the function W the graphon, we mean implicitly that a graphon is a tuple of an underlying probability space and a suitable function W . One useful aspect of a graphon W is that it naturally defines a family of random-graph models, sometimes called W -random graphs (see [[29](#), Chapter 11]). For every $n \in \mathbb{N}_{\geq 1}$, we denote by $\mathbb{G}_{W, n}$ a distribution on undirected graphs with vertex set $[n]$ that is induced by the following procedure for generating a random graph:

1. Draw a tuple $(x_1, \dots, x_n) \in X^n$ according to the product distribution ξ^n .
2. For all $i, j \in [n], i \neq j$, add the edge $\{i, j\}$ independently with probability $W(x_i, x_j)$.

Observe that $\mathbb{G}_{W, n}$ encompasses classical random-graph models, such as Erdős–Rényi random graphs and geometric random graphs.

Applying [Corollary 3.3](#) and using essentially the same arguments as in our proof sketch for [Theorem 1.1](#) yields the following result.

► **Theorem 1.4.** Let W be a graphon on the probability space $\mathcal{X} = (X, \mathcal{A}, \xi)$. Let $\gamma: \mathbb{N}_{\geq 1} \rightarrow \mathbb{R}_{\geq 0}$ such that $\gamma(n) \leq \gamma_0 n^{-\frac{1+\alpha}{2}}$ for some $\gamma_0 \in \mathbb{R}_{\geq 0}$ and $\alpha \in \mathbb{R}_{>0}$. For all $\beta \in [0, 1]$, $\varepsilon \in (0, 1]$, $\delta \in (0, 1]$, $n \geq (2\gamma_0^2 \varepsilon^{-2} \delta^{-1})^{\frac{1}{\alpha}}$, and $G \sim \mathbb{G}_{W,n}$, it holds that

$$\Pr[|Z_G(\gamma(n), \beta) - \mathbb{E}[Z_G(\gamma(n), \beta)]| \geq \varepsilon \mathbb{E}[Z_G(\gamma(n), \beta)]] \leq \delta. \quad \blacktriangleleft$$

► **Remark 1.5.** In fact, [Theorem 1.4](#) can easily be extended to an even more general setting, where we consider a sequence of probability spaces $(\mathcal{X}_n)_{n \in \mathbb{N}}$ and an associated sequence of graphons $(W_n)_{n \in \mathbb{N}}$ (i.e., each W_n is a graphon on \mathcal{X}_n). This makes the result for example applicable to popular models studied in network theory, such as hyperbolic random graphs [\[27\]](#) and geometric inhomogeneous random graphs [\[9\]](#). ◀

A surprising aspect of [Theorem 1.4](#) is that, even though $\mathbb{E}[Z_G(\gamma(n), \beta)] \geq 1 + n\gamma(n)$ diverges for $\gamma(n) \in \omega(n^{-1})$ as n increases, [Theorem 1.4](#) still ensures that the distribution of the partition functions gets more and more concentrated as long as $\gamma(n) \in o(n^{-\frac{1}{2}})$.

We derive [Theorem 1.1](#) as a special case of [Theorem 1.4](#). To see how this works, first observe that, for $\beta = 0$, it holds that $Z_G(\gamma(n), \beta)$ is the hard-core partition function of a graph G with parameter $\gamma \in \mathbb{R}_{\geq 0}$. Moreover, by setting $\mathcal{X} = (\mathbb{V}, \mathcal{B}, u_{\mathbb{V}})$, considering a graphon $W_{\phi}(x_1, x_2) = 1 - e^{-\phi(x_1, x_2)}$, $x_1, x_2 \in \mathbb{V}$, on \mathcal{X} , we obtain $\zeta_{\mathbb{V}, \phi}^{(n)} = \mathbb{G}_{W_{\phi}, n}$. This way of expressing $\zeta_{\mathbb{V}, \phi}^{(n)}$ establishes a connection between repulsive Gibbs point processes and hard-core models on graphon-based random graphs. Lastly, setting $\gamma_0 = \lambda\nu(\mathbb{V})$ and $\gamma(n) = \gamma_0 n^{-1}$ and applying [Theorem 1.4](#) yields the desired concentration result for hard-core partition functions on $\zeta_{\mathbb{V}, \phi}^{(n)}$.

1.6 Organization of the technical details

The technical details in the appendix are organized as follows. In [Section 2](#), we formally introduce the notion of antiferromagnetic spin systems and Gibbs point processes. Note that for the latter one, we base our definition on the notion of random counting measures. This is common in the theory of point processes but slightly differs from the definition given in the introduction. In [Section 3](#), we prove our concentration result for partition functions of antiferromagnetic spin systems on graphon-based random graphs. We continue by showing in [Section 4](#) how to efficiently approximate the partition functions of a Gibbs point process with arbitrary repulsive potentials via hard-core partition functions of random graphs based on a suitably constructed graphon. In [Section 5](#), we show how a similar approach can be used to obtain an approximate sampler for repulsive Gibbs point processes. Finally, in [Section 6](#) we give a high-probability bound on the connective constant of the graphs obtained from our reduction and derive corresponding strong spatial mixing results.

2 Preliminaries

We formally introduce the discrete antiferromagnetic spin systems we investigate, as well as Gibbs point processes.

2.1 Antiferromagnetic spin systems

For an undirected graph $G = (V, E)$ with vertices V and edges $E \subseteq \binom{V}{2}$, we denote by Σ_G the set of all functions $\sigma: V \rightarrow \{0, 1\}$. Without loss of generality, we are going to assume the canonical vertex set

$V = [n]$ for some $n \in \mathbb{N}$. A *two-state spin system* with parameters $\gamma, \beta_0, \beta_1 \in \mathbb{R}_{\geq 0}$ on G is a probability distribution $\mu_G^{(\gamma, \beta_0, \beta_1)}$ on Σ_G with

$$\mu_G^{(\gamma, \beta_0, \beta_1)}(\sigma) = \frac{\gamma^{|\sigma|_1} \beta_0^{m_G^{(0)}(\sigma)} \beta_1^{m_G^{(1)}(\sigma)}}{Z_G(\gamma, \beta_0, \beta_1)},$$

where $|\sigma|_1 = |\sigma^{-1}(1)|$ counts the number of vertices assigned that are to 1, $m_G^{(a)}(\sigma) = \sum_{\{i,j\} \in E} \mathbb{1}_{\sigma(i)=\sigma(j)=a}$ counts the number of edges with both endpoints assigned to $a \in \{0, 1\}$ and $Z_G(\gamma, \beta_0, \beta_1)$ is the normalizing constant

$$Z_G(\gamma, \beta_0, \beta_1) = \sum_{\sigma \in \Sigma_G} \gamma^{|\sigma|_1} \beta_0^{m_G^{(0)}(\sigma)} \beta_1^{m_G^{(1)}(\sigma)}.$$

Note that we implicitly assume $\beta_0 \neq 0$ or $\beta_1 \neq 0$, as $\mu_G^{(\gamma, \beta_0, \beta_1)}$ may not be defined otherwise.

Usually, $\mu_G^{(\gamma, \beta_0, \beta_1)}$ is referred to as the *Gibbs distribution* of the model and Z_G is called the *partition function*. Further, a two-state spin system is *antiferromagnetic* if $\beta_0 \beta_1 \leq 1$. For our concentration result, we focus on the setting where $\beta_0 = 1$. In this case, we omit β_0 completely and write $\beta = \beta_1 \in [0, 1]$ and denote the partition function by $Z_G(\gamma, \beta)$. Of special interest within this class of antiferromagnetic two-state spin systems in the *hard-core model*, which results from setting $\beta = 0$. In this case, we might just omit the edge interactions β completely and write $\mu_G^{(\gamma)}$ and $Z_G(\gamma)$. Note that this implies that only configurations $\sigma \in \Sigma_G$ for which $\sigma^{-1}(1)$ is an independent set in G can have non-zero probability. For us, this model is especially relevant, as we show that concentration of hard-core partition functions on random graphs can be used to derive randomized approximations for the partition function of repulsive Gibbs point processes, which are introduced in the next section.

2.2 Gibbs point processes

We introduce the notion of Gibbs point processes that is used throughout this paper. For a formal treatment, it is common to model point processes as random counting measures. Note that this is different from the simplified definition that we gave in the introduction. For a more detailed overview on the theory of point processes and specifically Gibbs point processes, see [25].

Let (\mathbb{X}, d) be a complete, separable metric space and let $\mathcal{B} = \mathcal{B}(\mathbb{X})$ be the Borel algebra of that space. Let ν be a locally finite reference measure on $(\mathbb{X}, \mathcal{B})$ such that all bounded measurable sets have finite measure. Denote by \mathcal{N} the set of all locally finite counting measures on $(\mathbb{X}, \mathcal{B})$. Formally, this is the set of all measures η on $(\mathbb{X}, \mathcal{B})$ with values in $\mathbb{N} \cup \{\infty\}$ such that $\nu(A) < \infty$ implies $\eta(A) < \infty$ for all $A \in \mathcal{B}$. For each $A \in \mathcal{B}$, define a map $N_A : \mathcal{N} \rightarrow \mathbb{N} \cup \{\infty\}$ with $\eta \mapsto \eta(A)$ and let \mathcal{R} be the sigma algebra on \mathcal{N} that is generated by the set of those maps $\{N_A \mid A \in \mathcal{B}\}$. A *point process* on \mathbb{X} is now a measurable map from some probability space to the measurable space $(\mathcal{N}, \mathcal{R})$. With some abuse of terminology, we call any probability distribution on $(\mathcal{N}, \mathcal{R})$ a point process, as we can only use the identity as measurable mapping from η to itself. Moreover, a point process is called *simple* if $N_x(\eta) \leq 1$ with probability 1, where we write N_x for $N_{\{x\}}$.

Note that every counting measure $\eta \in \mathcal{N}$ is associated with a multiset of points in \mathbb{X} . To see this, define $X_\eta = \{x \in \mathbb{X} \mid N_x(\eta) > 0\}$. Then η can be expressed as a weighted sum of Dirac measures

$$\eta = \sum_{x \in X_\eta} N_x(\eta) \delta_x.$$

In this sense, η is associated with a multiset of points $x \in X_\eta$, each occurring with finite multiplicity $N_x(\eta)$. We may use such a *point configuration* interchangeably with its corresponding counting measure.

An important example for point processes are Poisson point processes. A *Poisson point process* with intensity $\kappa \in \mathbb{R}_{\geq 0}$ on (\mathbb{X}, d) is uniquely defined by the following properties

- for all bounded measurable $A \subseteq \mathbb{X}$ it holds that N_A is Poisson distributed with intensity $\kappa\nu(A)$ and
- for all $m \in \mathbb{N}_{\geq 2}$ and disjoint measurable $A_1, \dots, A_m \subseteq \mathbb{X}$ it holds that N_{A_1}, \dots, N_{A_m} are independent.

Generally speaking, a *Gibbs point process* is a point process that is absolutely continuous with respect to a Poisson point process. For a bounded measurable $V \subseteq \mathbb{X}$ let \mathcal{N}_V denote the set of locally finite counting measures $\eta \in \mathcal{N}$ that satisfy $N_A(\eta) = 0$ for all measurable $A \subseteq \mathbb{X} \setminus V$. In this work we are interested in Gibbs point processes $P_V^{(\lambda, \phi)}$ on bounded measurable regions $V \subseteq \mathbb{X}$ that are parameterized by a *fugacity* parameter $\lambda \in \mathbb{R}_{\geq 0}$ and non-negative, symmetric, measurable *potential function* $\phi : \mathbb{X}^2 \rightarrow \mathbb{R}_{\geq 0} \cup \{\infty\}$. Formally, such a process $P_V^{(\lambda, \phi)}$ is defined by having a density with respect to a Poisson point process with intensity λ of the form

$$\frac{dP_V^{(\lambda, \phi)}}{dQ_\lambda}(\eta) = \frac{\mathbb{1}_{\eta \in \mathcal{N}_V} e^{-H(\eta)} e^{\lambda\nu(V)}}{\Xi_V(\lambda, \phi)}$$

where $H : \mathcal{N} \rightarrow \mathbb{R}_{\geq 0} \cup \{\infty\}$ is the *Hamiltonian* defined by

$$H(\eta) = \sum_{\{x, y\} \in \binom{X_\eta}{2}} N_x(\eta) N_y(\eta) \phi(x, y) + \sum_{x \in X_\eta} \frac{N_x(\eta)(N_x(\eta) - 1)}{2} \phi(x, x).$$

The normalizing constant $\Xi_V(\lambda, \phi)$ is usually called the (*grand-canonical*) *partition function* and can be written explicitly as

$$\begin{aligned} \Xi_V(\lambda, \phi) &= 1 + \sum_{k \in \mathbb{N}_{\geq 1}} \frac{\lambda^k}{k!} \int_{V^k} e^{-H(\delta_{x_1} + \dots + \delta_{x_k})} \nu^k(d\mathbf{x}) \\ &= 1 + \sum_{k \in \mathbb{N}_{\geq 1}} \frac{\lambda^k}{k!} \int_{V^k} \prod_{\{i, j\} \in \binom{[k]}{2}} e^{-\phi(x_i, x_j)} \nu^k(d\mathbf{x}). \end{aligned}$$

3 Concentration of partition functions of antiferromagnetic spin systems on graphon-based random graphs

The main tool we will use to derive our concentration bounds is the Efron–Stein inequality. For $N \in \mathbb{N}_{\geq 1}$ let $\{(\Omega_i, \mathcal{F}_i, \mu_i)\}_{i \in [N]}$ be a collection of probability spaces and let $f : \Omega \rightarrow \mathbb{R}$ be a measurable function on the product space $(\Omega, \mathcal{F}, \mu) = \bigotimes_{i \in [N]} (\Omega_i, \mathcal{F}_i, \mu_i)$. For each $i \in [N]$ define a function $\Delta_i^{(f)} : \Omega \times \Omega_i \rightarrow \mathbb{R}_{\geq 0}$, where, for every $\mathbf{x} = (x_1, \dots, x_N) \in \Omega$ and $y_i \in \Omega_i$, the value $\Delta_i^{(f)}(\mathbf{x}, y_i)$ is defined as the squared difference in f that is caused by replacing x_i in \mathbf{x} with y_i . Formally, this is $\Delta_i^{(f)}(\mathbf{x}, y_i) = (f(\mathbf{x}) - f(\mathbf{y}))^2$ where $\mathbf{y} = (x_1, \dots, x_{i-1}, y_i, x_{i+1}, \dots, x_N)$. The Efron–Stein inequality bounds the variance of f under μ based on the local squared deviations $\Delta_i^{(f)}(\mathbf{x}, y_i)$.

► **Theorem 3.1 (Efron–Stein inequality [14]).** Let $\{(\Omega_i, \mathcal{F}_i, \mu_i)\}_{i \in [N]}$ be probability spaces with product space $(\Omega, \mathcal{F}, \mu) = \bigotimes_{i \in [N]} (\Omega_i, \mathcal{F}_i, \mu_i)$. For every \mathcal{F} -measurable function $f : \Omega \rightarrow \mathbb{R}$ it holds that

$$\text{Var}_\mu[f] \leq \frac{1}{2} \sum_{i \in [N]} \mathbb{E}_{\mu \times \mu_i} \left[\Delta_i^{(f)} \right]. \quad \blacktriangleleft$$

► **Remark 3.2.** The Efron–Stein inequality is usually stated for functions of independent real-valued random variables. However, it extends to functions on products of arbitrary probability spaces. \blacktriangleleft

[Theorem 3.1](#) immediately gives a concentration result for f whenever $\frac{1}{2} \sum_{i \in [N]} \mathbb{E}_{\mu \times \mu_i} \left[\Delta_i^{(f)} \right]$ is of order of magnitude $\mathbb{E}_\mu[f]^2$ by using Chebyshev’s inequality. However, obtaining such a bound might turn out difficult, especially if $\mathbb{E}_\mu[f]$ is hard to compute explicitly. For our setting, we derive the following corollary of [Theorem 3.1](#).

► **Corollary 3.3.** [Corollary of the Efron–Stein inequality] Let $\{(\Omega_i, \mathcal{F}_i, \mu_i)\}_{i \in [N]}$ be probability spaces with product space $(\Omega, \mathcal{F}, \mu) = \bigotimes_{i \in [N]} (\Omega_i, \mathcal{F}_i, \mu_i)$, and let $f : \Omega \rightarrow \mathbb{R}$ be an \mathcal{F} -measurable function. Assume that there are $c_i \in \mathbb{R}_{\geq 0}$ for $i \in [N]$ such that $C := \sum_{i \in [N]} c_i^2 < 2$ and, for all $\mathbf{x} = (x_j)_{j \in [N]} \in \Omega$ and $\mathbf{y} = (y_j)_{j \in [N]} \in \Omega$ that disagree only at position i , it holds that

$$|f(\mathbf{x}) - f(\mathbf{y})| \leq c_i \cdot \min\{|f(\mathbf{x})|, |f(\mathbf{y})|\}.$$

Then, for all $\varepsilon \in \mathbb{R}_{>0}$, it holds that

$$\Pr[|f - \mathbb{E}_\mu[f]| \geq \varepsilon \mathbb{E}_\mu[f]] \leq \left(\frac{2}{2-C} - 1 \right) \frac{1}{\varepsilon^2}. \quad \blacktriangleleft$$

Proof. First, we observe that $|f(\mathbf{x}) - f(\mathbf{y})| \leq c_i \min\{|f(\mathbf{x})|, |f(\mathbf{y})|\} \leq c_i |f(\mathbf{x})|$ implies $\mathbb{E}_{\mu \times \mu_i} \left[\Delta_i^{(f)} \right] \leq c_i^2 \mathbb{E}_\mu[f^2]$ for all $i \in [N]$. Thus, by [Theorem 3.1](#), we have $\text{Var}_\mu[f] \leq \frac{C}{2} \mathbb{E}_\mu[f^2]$. Now, recall that by definition $\text{Var}_\mu[f] = \mathbb{E}_\mu[f^2] - \mathbb{E}_\mu[f]^2$, which implies $\mathbb{E}_\mu[f^2] - \mathbb{E}_\mu[f]^2 \leq \frac{C}{2} \mathbb{E}_\mu[f^2]$. Rearranging for $\mathbb{E}_\mu[f^2]$ and using the fact that $\frac{C}{2} < 1$ yields $\mathbb{E}_\mu[f^2] \leq \frac{2}{2-C} \mathbb{E}_\mu[f]^2$. Substituting this back into the definition of the variance, we obtain

$$\text{Var}_\mu[f] \leq \left(\frac{2}{2-C} - 1 \right) \mathbb{E}_\mu[f]^2.$$

The claim follows immediately by applying Chebyshev’s inequality. \blacksquare

► **Remark 3.4.** Usually, we want to characterize concentration asymptotically in N . In this setting, [Corollary 3.3](#) tells us that, if $c_i \in \mathcal{O}\left(N^{-\frac{1+\alpha}{2}}\right)$ for all $i \in [N]$ and some $\alpha > 0$, then, for all $\varepsilon \in \mathbb{R}_{>0}$ and $\delta \in (0, 1]$ such that $\varepsilon^2 \delta < 1$, it is sufficient to choose $N \in \Theta\left(\delta^{-\frac{1}{\alpha}} \varepsilon^{-\frac{2}{\alpha}}\right)$ to ensure

$$\Pr[|f - \mathbb{E}_\mu[f]| \geq \varepsilon \mathbb{E}_\mu[f]] \leq \delta. \quad \blacktriangleleft$$

We are now ready to use [Corollary 3.3](#) and derive a concentration result for the partition functions of antiferromagnetic two-state spin systems for graphon-based random graph model.

Let us recall the definition of graphons and graphon-based random graphs that we are using (see [29, Chapter 10 & 13]). Let $\mathcal{X} = (X, \mathcal{A}, \xi)$ be a probability space. A graphon on \mathcal{X} is a symmetric function $W : X^2 \rightarrow [0, 1]$ that is measurable with respect to the product algebra $\mathcal{A}^2 = \mathcal{A} \otimes \mathcal{A}$. For $n \in \mathbb{N}_{\geq 1}$ we denote by \mathcal{G}_n the set of all graphs on the canonical vertex set $[n] = \{1, \dots, n\}$. Note that each graph in \mathcal{G}_n

is fully characterized by its edge set E . For every $n \in \mathbb{N}_{\geq 1}$ the random graph model induced by a graphon W on a probability space (X, \mathcal{A}, ξ) is described by generating a random graph $G = ([n], E)$ by

- drawing a tuple $(x_1, \dots, x_n) \in X^n$ according to the product distribution ξ^n and
- adding the edge $\{i, j\}$ for all $i, j \in [n], i \neq j$ independently with probability $W(x_i, x_j)$.

Formally, this gives a probability distribution $\mathbb{G}_{W,n}$ on \mathcal{G}_n with

$$\mathbb{G}_{W,n}(G) = \int_{X^n} \left(\prod_{\{i,j\} \in E} W(x_i, x_j) \right) \cdot \left(\prod_{\{i,j\} \in \binom{[n]}{2} \setminus E} (1 - W(x_i, x_j)) \right) \xi^n(d\mathbf{x})$$

for all $G \in \mathcal{G}_n$, where $\mathbf{x} = (x_i)_{i \in [n]}$ inside the integral.

To apply [Corollary 3.3](#) to partition functions on random graphs from $\mathbb{G}_{W,n}$, we will need to bound how much the partition function changes when applying small modifications to the structure of a graph. More specifically, we want to get a bound on the relative change of the partition function, given that we

- add or remove a single edge, or
- add or remove a set of edges that are all incident to the same vertex.

The following two lemmas provide such bounds.

► **Lemma 3.5.** Let $G = (V, E)$ be an undirected graph and, for any $e \in E$ let $G' = (V, E \setminus \{e\})$. For all $\gamma \in \mathbb{R}_{\geq 0}$ and $\beta \in [0, 1]$ it holds that

$$0 \leq Z_{G'}(\gamma, \beta) - Z_G(\gamma, \beta) \leq \gamma^2 Z_G(\gamma, \beta)$$

and especially

$$|Z_{G'}(\gamma, \beta) - Z_G(\gamma, \beta)| \leq \gamma^2 \min\{Z_G(\gamma, \beta), Z_{G'}(\gamma, \beta)\}. \quad \blacktriangleleft$$

Proof. Without loss of generality, assume $V = [n]$ for some $n \in \mathbb{N}_{\geq 2}$ and let $e = \{i, j\}$ for $i, j \in [n]$. Note that $\Sigma_G = \Sigma_{G'}$, as their vertex sets are identical. Further, observe that, for all $\sigma \in \Sigma_G$, it holds that $m_G^{(1)}(\sigma) \geq m_{G'}^{(1)}(\sigma)$. Thus, we have $\beta^{m_G^{(1)}(\sigma)} \leq \beta^{m_{G'}^{(1)}(\sigma)}$ and $Z_G(\gamma, \beta) \leq Z_{G'}(\gamma, \beta)$, which proves

$$0 \leq Z_{G'}(\gamma, \beta) - Z_G(\gamma, \beta).$$

We proceed by rewriting the partition function of G' as

$$Z_{G'}(\gamma, \beta) = \sum_{\substack{\sigma \in \Sigma_{G'}: \\ \sigma(i)=0 \text{ or } \sigma(j)=0}} \gamma^{|\sigma|_1} \beta^{m_{G'}^{(1)}(\sigma)} + \sum_{\substack{\sigma \in \Sigma_{G'}: \\ \sigma(i)=\sigma(j)=1}} \gamma^{|\sigma|_1} \beta^{m_{G'}^{(1)}(\sigma)}.$$

Observe that

$$\sum_{\substack{\sigma \in \Sigma_{G'}: \\ \sigma(i)=0 \text{ or } \sigma(j)=0}} \gamma^{|\sigma|_1} \beta^{m_{G'}^{(1)}(\sigma)} = \sum_{\substack{\sigma \in \Sigma_G: \\ \sigma(i)=0 \text{ or } \sigma(j)=0}} \gamma^{|\sigma|_1} \beta^{m_G^{(1)}(\sigma)} \leq Z_G(\gamma, \beta).$$

For every $k \in [n]$, let $N_{G'}(k)$ denote the neighbors of vertex k in G' . We have

$$\sum_{\substack{\sigma \in \Sigma_{G'}: \\ \sigma(i)=\sigma(j)=1}} \gamma^{|\sigma|_1} \beta^{m_{G'}^{(1)}(\sigma)} = \sum_{\substack{\sigma \in \Sigma_{G'}: \\ \sigma(i)=\sigma(j)=0}} \gamma^{|\sigma|_1+2} \beta^{m_{G'}^{(1)}(\sigma)} \beta^{\sum_{k \in N_{G'}(i)} \sigma(k)} \beta^{\sum_{k \in N_{G'}(j)} \sigma(k)}$$

$$\begin{aligned}
&\leq \gamma^2 \sum_{\substack{\sigma \in \Sigma_{G'} \\ \sigma(i)=\sigma(j)=0}} \gamma^{|\sigma|_1} \beta^{m_{G'}^{(1)}(\sigma)} \\
&= \gamma^2 \sum_{\substack{\sigma \in \Sigma_G \\ \sigma(i)=\sigma(j)=0}} \gamma^{|\sigma|_1} \beta^{m_G^{(1)}(\sigma)} \\
&\leq \gamma^2 Z_G(\gamma, \beta).
\end{aligned}$$

We conclude that $Z_{G'}(\gamma, \beta) \leq (1 + \gamma^2)Z_G(\gamma, \beta)$ and thus

$$Z_{G'}(\gamma, \beta) - Z_G(\gamma, \beta) \leq \gamma^2 Z_G(\gamma, \beta).$$

The upper bound on $|Z_{G'}(\gamma, \beta) - Z_G(\gamma, \beta)|$ follows immediately. ■

► **Lemma 3.6.** Let $G = (V, E)$ be an undirected graph and without loss of generality assume $V = [n]$ for $n \in \mathbb{N}$. Let $E_H, E_{H'} \subseteq \{\{n+1, i\} \mid i \in [n]\}$, and set $H = ([n+1], E \cup E_H)$ and $H' = ([n+1], E \cup E_{H'})$. For all $\gamma \in \mathbb{R}_{\geq 0}$ and $\beta \in [0, 1]$ it holds that

$$|Z_H(\gamma, \beta) - Z_{H'}(\gamma, \beta)| \leq \gamma Z_G(\gamma, \beta) \leq \gamma \min\{Z_H(\gamma, \beta), Z_{H'}(\gamma, \beta)\}. \quad \blacktriangleleft$$

Proof. By [Lemma 3.5](#), we know that removing an edge from a graph doesn't decrease the partition function. Thus, $Z_H(\gamma, \beta)$ is maximized by choosing $E_H = \emptyset$ and minimized by choosing $E_H = \{\{n+1, i\} \mid i \in [n]\}$. Consequently, we have

$$Z_H(\gamma, \beta) \leq (1 + \gamma)Z_G(\gamma, \beta)$$

and

$$Z_H(\gamma, \beta) \geq Z_G(\gamma, \beta) + \gamma \geq Z_G(\gamma, \beta).$$

As the same holds for H' , we obtain

$$|Z_H(\gamma, \beta) - Z_{H'}(\gamma, \beta)| \leq \gamma Z_G(\gamma, \beta)$$

and the claim follows by noting that $Z_G(\gamma, \beta) \leq \min\{Z_H(\gamma, \beta), Z_{H'}(\gamma, \beta)\}$. ■

Based on [Lemmas 3.5](#) and [3.6](#), we use [Corollary 3.3](#) to prove the following statement.

► **Theorem 3.7.** Let W be a graphon on the probability space $\mathcal{X} = (X, \mathcal{A}, \xi)$. Let $\gamma : \mathbb{N}_{\geq 1} \rightarrow \mathbb{R}_{\geq 0}$ such that $\gamma(n) \leq \gamma_0 n^{-\frac{1+\alpha}{2}}$ for some $\gamma_0 \in \mathbb{R}_{\geq 0}$ and $\alpha \in \mathbb{R}_{>0}$. For all $\beta \in [0, 1]$, $\varepsilon \in \mathbb{R}_{>0}$, $n > \gamma_0^{\frac{2}{\alpha}}$ and $G \sim \mathbb{G}_{W,n}$ it holds that

$$\Pr\left[|Z_G(\gamma(n), \beta) - \mathbb{E}[Z_G(\gamma(n), \beta)]| \geq \varepsilon \mathbb{E}[Z_G(\gamma(n), \beta)]\right] \leq \frac{\gamma_0^2}{(n^\alpha - \gamma_0^2)\varepsilon^2}. \quad \blacktriangleleft$$

Proof. We aim for applying [Corollary 3.3](#) to prove our claim. To this end, for each $n \in \mathbb{N}_{\geq 1}$ we need to write the partition function $Z_G(\gamma(n), \beta)$ for $G \sim \mathbb{G}_{W,n}$ as a function on a product of σ -finite probability spaces. At first, an obvious choice seems to be \mathcal{X}^n together with $\binom{n}{2}$ additional binary random variables, one for each potential edge $\{i, j\} \in \binom{[n]}{2}$. However, note that the edges might not necessarily be independent, meaning that the resulting product distribution would not resemble $\mathbb{G}_{W,n}$. Instead, let $\mathcal{Y} = ([0, 1], \mathcal{B}([0, 1]), u)$, where $\mathcal{B}([0, 1])$ is the Borel algebra restricted to $[0, 1]$ and u is the uniform distribution on that interval. We consider the probability space $\mathcal{X}^n \otimes \mathcal{Y}^{\binom{n}{2}}$.

For $\mathbf{x} \in X^n$ and $\mathbf{y} \in [0, 1]^{\binom{n}{2}}$ let $\mathbf{x} \circ \mathbf{y} \in X^n \times [0, 1]^{\binom{n}{2}}$ denote the concatenation of \mathbf{x} and \mathbf{y} . We construct a measurable function $g : X^n \times [0, 1]^{\binom{n}{2}} \rightarrow \mathcal{G}_n$ by mapping every $\mathbf{z} = \mathbf{x} \circ \mathbf{y} \in X^n \times [0, 1]^{\binom{n}{2}}$ with $\mathbf{x} = (x_i)_{i \in [n]} \in X^n$ and $\mathbf{y} = (y_{i,j})_{1 \leq i < j \leq n} \in [0, 1]^{\binom{n}{2}}$ to $g(\mathbf{z}) = ([n], E)$ such that, for all $i < j$, it holds that $\{i, j\} \in E$ if and only if $W(x_i, x_j) \geq y_{i,j}$. Simple calculations show that, for $\mathbf{z} \sim \xi^n \times u^{\binom{n}{2}}$, it holds that $g(\mathbf{z}) \sim \mathbb{G}_{W,n}$. Now, let $f : X^n \times [0, 1]^{\binom{n}{2}} \rightarrow \mathbb{R}$ with $\mathbf{z} \mapsto Z_{g(\mathbf{z})}(\gamma(n), \beta)$. In order to apply [Corollary 3.3](#), we need to bound the relative change of $f(\mathbf{z})$ if we change one component of \mathbf{z} . Let $\mathbf{x}' = (x_1, \dots, x_{i-1}, x'_i, x_{i+1}, \dots, x_n) \in X^n$ for any $i \in [n]$. Then $g(\mathbf{x}' \circ \mathbf{y})$ can only differ from $g(\mathbf{z})$ on edges that are incident to vertex i . Thus, by [Lemma 3.6](#), it holds that

$$|f(\mathbf{z}) - f(\mathbf{x}' \circ \mathbf{y})| \leq \gamma(n) \min\{f(\mathbf{z}), f(\mathbf{x}' \circ \mathbf{y})\}.$$

Now, let $\mathbf{y}' = (y'_{i,j})_{1 \leq i < j \leq n} \in [0, 1]^{\binom{n}{2}}$ such that $y'_{i,j} = y_{i,j}$ except for one pair $1 \leq i < j \leq n$. Note that $g(\mathbf{z})$ and $g(\mathbf{x} \circ \mathbf{y}')$ differ by at most one edge. By [Lemma 3.5](#), we have

$$|f(\mathbf{z}) - f(\mathbf{x} \circ \mathbf{y}')| \leq \gamma(n)^2 \min\{f(\mathbf{z}), f(\mathbf{x} \circ \mathbf{y}')\}.$$

Furthermore, note that for $\gamma(n) \leq \gamma_0 n^{-\frac{1+\alpha}{2}}$ and $n > \gamma_0^{\frac{2}{\alpha}}$ it holds that

$$C = n\gamma(n)^2 + \binom{n}{2}\gamma(n)^4 \leq \gamma_0^2 n^{-\alpha} + \gamma_0^4 n^{-2\alpha} \leq 2\gamma_0^2 n^{-\alpha} < 2.$$

Thus, by [Corollary 3.3](#) we obtain

$$\begin{aligned} \Pr[|Z_G(\gamma(n), \beta) - \mathbb{E}[Z_G(\gamma(n), \beta)]| \geq \varepsilon \mathbb{E}[Z_G(\gamma(n), \beta)]] &\leq \left(\frac{2}{2-C} - 1\right) \frac{1}{\varepsilon^2} \\ &\leq \left(\frac{1}{1-\gamma_0^2 n^{-\alpha}} - 1\right) \frac{1}{\varepsilon^2} \\ &= \frac{\gamma_0^2}{(n^\alpha - \gamma_0^2)\varepsilon^2}, \end{aligned}$$

which concludes the proof. ■

[Theorem 1.4](#) follows immediately from [Theorem 3.7](#).

► **Theorem 1.4.** Let W be a graphon on the probability space $\mathcal{X} = (X, \mathcal{A}, \xi)$. Let $\gamma : \mathbb{N}_{\geq 1} \rightarrow \mathbb{R}_{\geq 0}$ such that $\gamma(n) \leq \gamma_0 n^{-\frac{1+\alpha}{2}}$ for some $\gamma_0 \in \mathbb{R}_{\geq 0}$ and $\alpha \in \mathbb{R}_{>0}$. For all $\beta \in [0, 1]$, $\varepsilon \in (0, 1]$, $\delta \in (0, 1]$, $n \geq (2\gamma_0^2 \varepsilon^{-2} \delta^{-1})^{\frac{1}{\alpha}}$, and $G \sim \mathbb{G}_{W,n}$, it holds that

$$\Pr[|Z_G(\gamma(n), \beta) - \mathbb{E}[Z_G(\gamma(n), \beta)]| \geq \varepsilon \mathbb{E}[Z_G(\gamma(n), \beta)]] \leq \delta. \quad \blacktriangleleft$$

Proof. For $\varepsilon \leq 1$ and $\delta \leq 1$ it holds that $n \geq (2\gamma_0^2 \varepsilon^{-2} \delta^{-1})^{\frac{1}{\alpha}} > \gamma_0^{\frac{2}{\alpha}}$. Applying [Theorem 3.7](#) yields

$$\begin{aligned} \Pr[|Z_G(\gamma(n), \beta) - \mathbb{E}[Z_G(\gamma(n), \beta)]| \geq \varepsilon \mathbb{E}[Z_G(\gamma(n), \beta)]] &\leq \frac{1}{(2\varepsilon^{-2}\delta^{-1} - 1)\varepsilon^2} \\ &= \frac{\varepsilon^2 \delta}{(2 - \varepsilon^2 \delta)\varepsilon^2} \\ &= \frac{\delta}{2 - \varepsilon^2 \delta} \end{aligned}$$

$\leq \delta$. ■

4 Application to repulsive Gibbs point processes

We use our concentration results for antiferromagnetic spin systems to relate repulsive Gibbs point processes to a hard-core model on carefully constructed classes of random graphs. To this end, let (\mathbb{X}, d) be a complete, separable metric space, let $\mathcal{B} = \mathcal{B}(\mathbb{X})$ be the Borel algebra and let ν be a locally finite reference measure on $(\mathbb{X}, \mathcal{B})$. For every bounded and measurable $\mathbb{V} \subseteq \mathbb{X}$ we define a probability space $\mathcal{X}_{\mathbb{V}} = (\mathbb{V}, \mathcal{B}_{\mathbb{V}}, u_{\mathbb{V}})$, where $\mathcal{B}_{\mathbb{V}}$ denotes the restriction of \mathcal{B} to \mathbb{V} and $u_{\mathbb{V}}$ is the probability measure on $(\mathbb{V}, \mathcal{B}_{\mathbb{V}})$ that is defined via the constant density $\frac{1}{\nu(\mathbb{V})}$ with respect to ν restricted to \mathbb{V} . For every symmetric, repulsive and measurable pair potential function $\phi : \mathbb{X}^2 \rightarrow \mathbb{R}_{\geq 0} \cup \{\infty\}$ and all $n \in \mathbb{N}_{\geq 1}$. Define $W_{\phi} : \mathbb{V}^2 \rightarrow [0, 1]$ with $W_{\phi}(x, y) = 1 - e^{-\phi(x, y)}$ and observe that W_{ϕ} is a graphon on $\mathcal{X}_{\mathbb{V}}$. We proceed by considering the random graph model $\zeta_{\mathbb{V}, \phi}^{(n)} = \mathbb{G}_{W_{\phi}, n}$.

The following lemma relates the expected hard-core partition function on $\zeta_{\mathbb{V}, \phi}^{(n)}$ with the partition function of the continuous Gibbs point process $\Xi_{\mathbb{V}}(\lambda, \phi)$.

► **Lemma 4.1.** Let (\mathbb{X}, d) be a complete separable metric space, let $\mathcal{B} = \mathcal{B}(\mathbb{X})$ be the Borel algebra and let ν be a locally finite reference measure on $(\mathbb{X}, \mathcal{B})$. Let $\mathbb{V} \subseteq \mathbb{X}$ be bounded and measurable, let $\lambda \in \mathbb{R}_{\geq 0}$ and let $\phi : \mathbb{X}^2 \rightarrow \mathbb{R}_{\geq 0} \cup \{\infty\}$ be a symmetric repulsive potential. For all $\varepsilon \in \mathbb{R}_{> 0}$ and $n \geq 2\varepsilon^{-1} \max\{e^6 \lambda^2 \nu(\mathbb{V})^2, \ln(2\varepsilon^{-1})^2\}$ it holds that

$$(1 - \varepsilon)\Xi_{\mathbb{V}}(\lambda, \phi) \leq \mathbb{E}_{G \sim \zeta_{\mathbb{V}, \phi}^{(n)}} \left[Z_G \left(\frac{\lambda \nu(\mathbb{V})}{n} \right) \right] \leq \Xi_{\mathbb{V}}(\lambda, \phi). \quad \blacktriangleleft$$

Proof. We start by rewriting the hard-core partition function as

$$Z_G \left(\frac{\lambda \nu(\mathbb{V})}{n} \right) = 1 + \sum_{k=1}^n \lambda^k \frac{\nu(\mathbb{V})^k}{n^k} \sum_{S \in \binom{[n]}{k}} \prod_{\{i, j\} \in \binom{S}{2}} \mathbb{1}_{\{i, j\} \notin E}.$$

Thus, by linearity of expectation we have

$$\begin{aligned} \mathbb{E}_{G \sim \zeta_{\mathbb{V}, \phi}^{(n)}} \left[Z_G \left(\frac{\lambda \nu(\mathbb{V})}{n} \right) \right] &= 1 + \sum_{k=1}^n \lambda^k \frac{\nu(\mathbb{V})^k}{n^k} \sum_{S \in \binom{[n]}{k}} \mathbb{E}_{G \sim \zeta_{\mathbb{V}, \phi}^{(n)}} \left[\prod_{\{i, j\} \in \binom{S}{2}} \mathbb{1}_{\{i, j\} \notin E} \right] \\ &= 1 + \sum_{k=1}^n \lambda^k \frac{\nu(\mathbb{V})^k}{n^k} \sum_{S \in \binom{[n]}{k}} \Pr \left[\bigwedge_{\{i, j\} \in \binom{S}{2}} \{i, j\} \notin E \right]. \end{aligned}$$

Next, observe that for all $S \in \binom{[n]}{k}$ with $|S| = k$

$$\Pr \left[\bigwedge_{\{i, j\} \in \binom{S}{2}} \{i, j\} \notin E \right] = \int_{\mathbb{V}^n} \prod_{\{i, j\} \in \binom{S}{2}} (1 - W_{\phi}(x_i, x_j)) u_{\mathbb{V}}^n(\mathbf{dx})$$

$$\begin{aligned}
&= \int_{\mathbb{V}^n} \prod_{\{i,j\} \in \binom{S}{2}} e^{-\phi(x_i, x_j)} u_{\mathbb{V}}^n(\mathbf{d}\mathbf{x}) \\
&= \frac{1}{v(\mathbb{V})^n} \int_{\mathbb{V}^n} \prod_{\{i,j\} \in \binom{S}{2}} e^{-\phi(x_i, x_j)} v^n(\mathbf{d}\mathbf{x}) \\
&= \frac{v(\mathbb{V})^{n-k}}{v(\mathbb{V})^n} \int_{\mathbb{V}^k} \prod_{\{i,j\} \in \binom{[k]}{2}} e^{-\phi(x_i, x_j)} v^k(\mathbf{d}\mathbf{x}) \\
&= \frac{1}{v(\mathbb{V})^k} \int_{\mathbb{V}^k} \prod_{\{i,j\} \in \binom{[k]}{2}} e^{-\phi(x_i, x_j)} v^k(\mathbf{d}\mathbf{x}).
\end{aligned}$$

This yields

$$\begin{aligned}
\mathbb{E}_{G \sim \zeta_{\mathbb{V}, \phi}^{(n)}} \left[Z_G \left(\frac{\lambda v(\mathbb{V})}{n} \right) \right] &= 1 + \sum_{k=1}^n \lambda^k \frac{1}{n^k} \sum_{S \in \binom{[n]}{k}} \int_{\mathbb{V}^k} \prod_{\{i,j\} \in \binom{[k]}{2}} e^{-\phi(x_i, x_j)} v^k(\mathbf{d}\mathbf{x}) \\
&= 1 + \sum_{k=1}^n \lambda^k \frac{\binom{n}{k}}{n^k} \int_{\mathbb{V}^k} \prod_{\{i,j\} \in \binom{[k]}{2}} e^{-\phi(x_i, x_j)} v^k(\mathbf{d}\mathbf{x}) \\
&= 1 + \sum_{k=1}^n \frac{\lambda^k}{k!} \prod_{i=0}^{k-1} \left(1 - \frac{i}{n} \right) \int_{\mathbb{V}^k} \prod_{\{i,j\} \in \binom{[k]}{2}} e^{-\phi(x_i, x_j)} v^k(\mathbf{d}\mathbf{x}),
\end{aligned}$$

from which the upper bound

$$\mathbb{E}_{G \sim \zeta_{\mathbb{V}, \phi}^{(n)}} \left[Z_G \left(\frac{\lambda v(\mathbb{V})}{n} \right) \right] \leq \Xi_{\mathbb{V}}(\lambda, \phi)$$

follows immediately.

For the lower bound set

$$S_m = 1 + \sum_{k=1}^m \frac{\lambda^k}{k!} \int_{\mathbb{V}^k} \prod_{\{i,j\} \in \binom{[k]}{2}} e^{-\phi(x_i, x_j)} v^k(\mathbf{d}\mathbf{x})$$

for any $1 \leq m \leq n$. Observe that

$$\mathbb{E}_{G \sim \zeta_{\mathbb{V}, \phi}^{(n)}} \left[Z_G \left(\frac{\lambda v(\mathbb{V})}{n} \right) \right] \geq \left(1 - \frac{m}{n} \right)^m S_m.$$

Thus, for $n \geq 2\varepsilon^{-1}m^2$ Bernoulli's inequality yields

$$\mathbb{E}_{G \sim \zeta_{\mathbb{V}, \phi}^{(n)}} \left[Z_G \left(\frac{\lambda v(\mathbb{V})}{n} \right) \right] \geq \left(1 - \frac{m^2}{n} \right) S_m \geq \left(1 - \frac{\varepsilon}{2} \right) S_m.$$

Furthermore, note that

$$\Xi_{\mathbb{V}}(\lambda, \phi) - S_m = \sum_{k=m+1}^{\infty} \frac{\lambda^k}{k!} \int_{\mathbb{V}^k} \prod_{\{i,j\} \in \binom{[k]}{2}} e^{-\phi(x_i, x_j)} v^k(\mathbf{d}\mathbf{x})$$

$$\leq \sum_{k=m+1}^{\infty} \frac{\lambda^k \nu(\mathbb{V})^k}{k!},$$

where the last inequality comes from the fact that ϕ is non-negative. Next, observe that this is equal to the error of the Taylor expansion of $e^{\lambda \nu(\mathbb{V})}$ around 0, truncated after m terms. Thus, by Lagrange's remainder formula, we obtain

$$\Xi_{\mathbb{V}}(\lambda, \phi) - S_m \leq \frac{e^{\lambda \nu(\mathbb{V})}}{(m+1)!} (\lambda \nu(\mathbb{V}))^{m+1}.$$

Choosing $m \geq \max\{e^3 \lambda \nu(\mathbb{V}), \ln(2\varepsilon^{-1})\}$ and using the fact that $(m+1)! > \left(\frac{m+1}{e}\right)^{m+1}$ yields

$$\Xi_{\mathbb{V}}(\lambda, \phi) - S_m \leq \left(\frac{e^2 \lambda \nu(\mathbb{V})}{m+1}\right)^{m+1} \leq e^{-(m+1)} \leq \frac{\varepsilon}{2}.$$

As $\Xi_{\mathbb{V}}(\lambda, \phi) \geq 1$, we get

$$S_m \geq \Xi_{\mathbb{V}}(\lambda, \phi) - \frac{\varepsilon}{2} \geq \left(1 - \frac{\varepsilon}{2}\right) \Xi_{\mathbb{V}}(\lambda, \phi).$$

For $n \geq 2\varepsilon^{-1}m^2 = 2\varepsilon^{-1} \max\{e^6 \lambda^2 \nu(\mathbb{V})^2, \ln(2\varepsilon^{-1})^2\}$ we obtain

$$\mathbb{E}_{G \sim \zeta_{\mathbb{V}, \phi}^{(n)}} \left[Z_G \left(\frac{\lambda \nu(\mathbb{V})}{n} \right) \right] \geq \left(1 - \frac{m}{n}\right)^m S_m \geq \left(1 - \frac{\varepsilon}{2}\right)^2 \Xi_{\mathbb{V}}(\lambda, \phi) \geq (1 - \varepsilon) \Xi_{\mathbb{V}}(\lambda, \phi),$$

which proves the claim. ■

► **Theorem 1.1.** Let (\mathbb{X}, d) be a complete, separable metric space, let $\mathcal{B} = \mathcal{B}(\mathbb{X})$ be the Borel algebra, and let ν be a locally finite reference measure on $(\mathbb{X}, \mathcal{B})$. Let $\mathbb{V} \subseteq \mathbb{X}$ be bounded and measurable, let $\lambda \in \mathbb{R}_{\geq 0}$, and let $\phi: \mathbb{X}^2 \rightarrow \mathbb{R}_{\geq 0} \cup \{\infty\}$ be a symmetric repulsive potential. For all $\varepsilon \in (0, 1]$, $\delta \in (0, 1]$ and $n \geq 4\varepsilon^{-2}\delta^{-1} \max\{e^6 \lambda^2 \nu(\mathbb{V})^2, \ln(4\varepsilon^{-1})^2\}$, it holds that, for $G \sim \zeta_{\mathbb{V}, \phi}^{(n)}$,

$$\Pr \left[\left| Z_G \left(\frac{\lambda \nu(\mathbb{V})}{n} \right) - \Xi_{\mathbb{V}}(\lambda, \phi) \right| \geq \varepsilon \Xi_{\mathbb{V}}(\lambda, \phi) \right] \leq \delta. \quad \blacktriangleleft$$

Proof. By setting $\alpha = 1$ and $\gamma_0 = \lambda \nu(\mathbb{V})$ and using the fact that

$$n \geq 4\varepsilon^{-2}\delta^{-1} \max\{e^6 \lambda^2 \nu(\mathbb{V})^2, \ln(4\varepsilon^{-1})^2\} \geq \left(2\gamma_0^2 \left(\frac{\varepsilon}{2}\right)^{-2} \delta^{-1}\right)^{\frac{1}{\alpha}}$$

Theorem 1.4 yields

$$\Pr \left[\left| Z_G \left(\frac{\lambda \nu(\mathbb{V})}{n} \right) - \mathbb{E} \left[Z_G \left(\frac{\lambda \nu(\mathbb{V})}{n} \right) \right] \right| \geq \frac{\varepsilon}{2} \mathbb{E} \left[Z_G \left(\frac{\lambda \nu(\mathbb{V})}{n} \right) \right] \right] \leq \delta.$$

Furthermore, by Lemma 4.1 we know that for

$$n \geq 4\varepsilon^{-2}\delta^{-1} \max\{e^6 \lambda^2 \nu(\mathbb{V})^2, \ln(4\varepsilon^{-1})^2\} \geq 2 \left(\frac{\varepsilon}{2}\right)^{-1} \max\left\{e^6 \lambda^2 \nu(\mathbb{V})^2, \ln\left(2 \left(\frac{\varepsilon}{2}\right)^{-1}\right)^2\right\}$$

it holds that

$$\left(1 - \frac{\varepsilon}{2}\right) \Xi_{\mathbb{V}}(\lambda, \phi) \leq \mathbb{E}_{G \sim \zeta_{\mathbb{V}, \phi}^{(n)}} \left[Z_G \left(\frac{\lambda \nu(\mathbb{V})}{n} \right) \right] \leq \Xi_{\mathbb{V}}(\lambda, \phi).$$

Thus, we have

$$\left(1 + \frac{\varepsilon}{2}\right) \mathbb{E} \left[Z_G \left(\frac{\lambda \nu(\mathbb{V})}{n} \right) \right] \leq \left(1 + \frac{\varepsilon}{2}\right) \Xi_{\mathbb{V}}(\lambda, \phi) \leq (1 + \varepsilon) \Xi_{\mathbb{V}}(\lambda, \phi)$$

and similarly

$$\left(1 - \frac{\varepsilon}{2}\right) \mathbb{E} \left[Z_G \left(\frac{\lambda \nu(\mathbb{V})}{n} \right) \right] \geq \left(1 - \frac{\varepsilon}{2}\right)^2 \Xi_{\mathbb{V}}(\lambda, \phi) \geq (1 - \varepsilon) \Xi_{\mathbb{V}}(\lambda, \phi).$$

We obtain

$$\Pr \left[\left| Z_G \left(\frac{\lambda \nu(\mathbb{V})}{n} \right) - \Xi_{\mathbb{V}}(\lambda, \phi) \right| \geq \varepsilon \Xi_{\mathbb{V}}(\lambda, \phi) \right] \leq \delta,$$

which proves the claim. ■

4.1 Approximating the partition function

One of the main applications of [Theorem 1.1](#) is that it yields a rather simple randomized procedure for approximating $\Xi_{\mathbb{V}}(\lambda, \phi)$. The rough idea is as follows:

1. For $n \in \mathbb{N}$ sufficiently large, sample a graph G from $\zeta_{\mathbb{V}, \phi}^{(n)}$.
2. Approximate $Z_G \left(\frac{\lambda \nu(\mathbb{V})}{n} \right)$ and use the result as an approximation for $\Xi_{\mathbb{V}}(\lambda, \phi)$.

We are especially interested in obtaining an algorithm that is asymptotically efficient in the volume $\nu(\mathbb{V})$, as this gives a natural way to parameterize the algorithmic problem. More specifically, we want to characterize the regime of the fugacity λ in terms of the potential ϕ for which we can get a randomized ε -approximation of $\Xi_{\mathbb{V}}(\lambda, \phi)$ in time polynomial in $\nu(\mathbb{V})$ and $\frac{1}{\varepsilon}$. We characterize this fugacity regime in terms of the temperedness constant

$$C_{\phi} = \operatorname{ess\,sup}_{x_1 \in \mathbb{X}} \int_{\mathbb{X}} \left| 1 - e^{-\phi(x_1, x_2)} \right| \nu(dx_2),$$

where $\operatorname{ess\,sup}$ denotes the essential supremum (i.e., an upper bound that holds almost everywhere).

In order to ensure that the approximation algorithm runs efficiently in $\nu(\mathbb{V})$, two ingredients are important. First, we need to bound how large n needs to be chosen to ensure that $Z_G \left(\frac{\lambda \nu(\mathbb{V})}{n} \right)$ is close to $\Xi_{\mathbb{V}}(\lambda, \phi)$ with high probability. Second, we need to ensure that $Z_G \left(\frac{\lambda \nu(\mathbb{V})}{n} \right)$ can be approximated in time polynomial in $\nu(\mathbb{V})$. Obviously, both requirements are satisfied if $n \in \operatorname{poly}(\nu(\mathbb{V}))$ is sufficient and if $Z_G \left(\frac{\lambda \nu(\mathbb{V})}{n} \right)$ can be approximated in time $\operatorname{poly}(n)$. To tackle the first part, [Theorem 1.1](#) gives a useful tool. For the second part, we will use some well known results on approximating the hard-core partition function.

► **Theorem 4.2** ([\[44, Corollary 8.4\]](#) and [\[3, Theorem 1\]](#)). Let $G = (V, E)$ be an undirected graph with maximum vertex degree bounded by $d_G \in \mathbb{N}_{\geq 2}$ and let $\gamma \in \mathbb{R}_{\geq 0}$ with

$$\gamma < \gamma_c(d_G) = \frac{(d_G - 1)^{d_G - 1}}{(d_G - 2)^{d_G}}.$$

Then, for all $\varepsilon \in (0, 1]$, there is a randomized ε -approximation algorithm for the hard-core partition function $Z_G(\gamma)$ with running time $\tilde{O}(|V|^2 \varepsilon^{-2})$. \blacktriangleleft

► **Remark 4.3.** In [44] the result above is only stated for $\gamma < \frac{2}{d_G}$ as an older mixing time result for Glauber dynamics from [46] is used. Combining their approach with the more recent mixing time bound in [3] gives the desired bound of $\gamma < \gamma_c(d_G)$. \blacktriangleleft

Thus, arguing that Z_G for $G \sim \zeta_{\mathbb{V}, \phi}^{(n)}$ can be approximated in time $\text{poly}(n)$ boils down to obtaining a probabilistic upper bound on d_G . We use the following simple lemma.

► **Lemma 4.4.** Let (\mathbb{X}, d) be a complete, separable metric space, let $\mathcal{B} = \mathcal{B}(\mathbb{X})$ be the Borel algebra and let ν be a locally finite reference measure on $(\mathbb{X}, \mathcal{B})$. Let $\mathbb{V} \subseteq \mathbb{X}$ be bounded and measurable, let $\lambda \in \mathbb{R}_{\geq 0}$ and let $\phi : \mathbb{X}^2 \rightarrow \mathbb{R}_{\geq 0} \cup \{\infty\}$ be a symmetric repulsive potential. Assume $C_\phi > 0$. For $\alpha \in \mathbb{R}_{> 0}$, $q \in (0, 1]$, $n \geq 3 \max\{\alpha^{-1}, \alpha^{-2}\} \ln(q^{-1}) C_\phi^{-1} \nu(\mathbb{V}) + 1$ and $G \sim \zeta_{\mathbb{V}, \phi}^{(n)}$ it holds that

$$\Pr \left[d_G \geq (1 + \alpha) \frac{n-1}{\nu(\mathbb{V})} C_\phi \right] \leq qn. \quad \blacktriangleleft$$

Proof. By union bound, it is sufficient to argue that, for each $i \in [n]$ it holds that

$$\Pr \left[d_G(i) \geq (1 + \alpha) \frac{n-1}{\nu(\mathbb{V})} C_\phi \right] \leq q,$$

where $d_G(i)$ denotes the degree of vertex $i \in [n]$ in G . Now, observe that the random variables $d_G(i)$ for $i \in [n]$ are identically distributed. Thus, we can focus on $d_G(n)$ for ease of notation. By definition, it holds for $k \in [n-1] \cup \{0\}$ that

$$\begin{aligned} \Pr[d_G(n) = k] &= \sum_{S \in \binom{[n-1]}{k}} \int_{\mathbb{V}^n} \left(\prod_{i \in S} W_\phi(x_n, x_i) \right) \cdot \left(\prod_{i \in [n-1] \setminus S} (1 - W_\phi(x_n, x_i)) \right) u_{\mathbb{V}}^n(\mathbf{x}) \\ &= \int_{\mathbb{V}} \sum_{S \in \binom{[n-1]}{k}} \left(\prod_{i \in S} \int_{\mathbb{V}} W_\phi(x_n, x_i) u_{\mathbb{V}}(\mathbf{d}x_i) \right) \cdot \left(\prod_{i \in [n-1] \setminus S} \int_{\mathbb{V}} 1 - W_\phi(x_n, x_i) u_{\mathbb{V}}(\mathbf{d}x_i) \right) u_{\mathbb{V}}(\mathbf{d}x_n) \\ &= \int_{\mathbb{V}} \binom{n-1}{k} \left(\int_{\mathbb{V}} W_\phi(x_1, x_2) u_{\mathbb{V}}(\mathbf{d}x_2) \right)^k \left(1 - \int_{\mathbb{V}} W_\phi(x_1, x_2) u_{\mathbb{V}}(\mathbf{d}x_2) \right)^{n-1-k} u_{\mathbb{V}}(\mathbf{d}x_1). \end{aligned}$$

For every $x_1 \in \mathbb{V}$, let B_{x_1} be a binomial random variable with $n-1$ trials and with success probability $\int_{\mathbb{V}} W_\phi(x_1, x_2) u_{\mathbb{V}}(\mathbf{d}x_2)$. We obtain

$$\Pr[d_G(n) = k] = \int_{\mathbb{V}} \Pr[B_{x_1} = k] u_{\mathbb{V}}(\mathbf{d}x_1),$$

which implies for all $a \in [0, n-1]$

$$\Pr[d_G(n) \geq a] = \sum_{k=\lceil a \rceil}^{n-1} \int_{\mathbb{V}} \Pr[B_{x_1} = k] u_{\mathbb{V}}(\mathbf{d}x_1)$$

$$\begin{aligned}
&= \int_{\mathbb{V}} \sum_{k=\lceil a \rceil}^{n-1} \Pr[B_{x_1} = k] u_{\mathbb{V}}(dx_1) \\
&= \int_{\mathbb{V}} \Pr[B_{x_1} \geq a] u_{\mathbb{V}}(dx_1).
\end{aligned}$$

Next, let B be a binomial random variable with $n - 1$ trials and success probability $\frac{C_\phi}{v(\mathbb{V})}$. Observe that, by the definition of C_ϕ , it holds for ν -almost all $x_1 \in \mathbb{V}$ that $\int_{\mathbb{V}} W_\phi(x_1, x_2) u_{\mathbb{V}}(dx_2) \leq \frac{C_\phi}{v(\mathbb{V})}$. Thus, we have that B stochastically dominates B_{x_1} for $u_{\mathbb{V}}$ -almost all $x_1 \in \mathbb{V}$. Consequently, we obtain

$$\Pr[d_G(n) \geq a] \leq \int_{\mathbb{V}} \Pr[B \geq a] u_{\mathbb{V}}(dx_1) = \Pr[B \geq a].$$

Observing that $E[B] = \frac{n-1}{v(\mathbb{V})}C_\phi$ and applying Chernoff bound yields

$$\Pr\left[d_G(n) \geq (1 + \alpha) \frac{n-1}{v(\mathbb{V})}C_\phi\right] \leq e^{-\frac{\min\{\alpha, \alpha^2\}C_\phi(n-1)}{3v(\mathbb{V})}}.$$

Setting $n \geq 3 \max\{\alpha^{-1}, \alpha^{-2}\} \ln(q^{-1})C_\phi^{-1}v(\mathbb{V}) + 1$ we have $\Pr\left[d_G(n) \geq (1 + \alpha) \frac{n-1}{v(\mathbb{V})}C_\phi\right] \leq q$, which proves the claim. \blacksquare

Combining [Theorem 1.1](#), [Lemma 4.4](#), and [Theorem 4.2](#), we obtain the following algorithmic result.

► **Theorem 1.2.** Let (\mathbb{X}, d) be a complete, separable metric space, let $\mathcal{B} = \mathcal{B}(\mathbb{X})$ be the Borel algebra, and let ν be a locally finite reference measure on $(\mathbb{X}, \mathcal{B})$. Let $\mathbb{V} \subseteq \mathbb{X}$ be bounded and measurable, let $\lambda \in \mathbb{R}_{\geq 0}$, and let $\phi: \mathbb{X}^2 \rightarrow \mathbb{R}_{\geq 0} \cup \{\infty\}$ be a symmetric repulsive potential. Assume there is a sampler for $\zeta_{\mathbb{V}, \phi}^{(n)}$ with running time $t_{\mathbb{V}, \phi}(n)$.

If $\lambda < \frac{\epsilon}{C_\phi}$, then, for all $\epsilon \in (0, 1]$, there is a randomized ϵ -approximation algorithm for $\Xi_{\mathbb{V}}(\lambda, \phi)$ with running time in $\tilde{O}(v(\mathbb{V})^4 \epsilon^{-6}) + t_{\mathbb{V}, \phi}(\tilde{O}(v(\mathbb{V})^2 \epsilon^{-2}))$. \blacktriangleleft

Proof. We start by giving a more precise outline of the algorithmic idea. To this end, we define

$$N = \max \left\{ \begin{array}{l} 324\epsilon^{-2} \max\{e^6 \lambda^2 v(\mathbb{V})^2, \ln(4\epsilon^{-1})^2\}, \\ 24 \max\left\{ \frac{1}{e^{-\lambda C_\phi}}, \frac{\lambda C_\phi}{(e^{-\lambda C_\phi})^2} \right\} \lambda v(\mathbb{V}) \ln \left(24 \max\left\{ \frac{1}{e^{-\lambda C_\phi}}, \frac{\lambda C_\phi}{(e^{-\lambda C_\phi})^2} \right\} \lambda v(\mathbb{V}) \right)^2 \end{array} \right\}.$$

We now use the following procedure to approximate $\Xi_{\mathbb{V}}(\lambda, \phi)$:

1. Choose some integer $n \geq N$.
2. Draw a graph G from $\zeta_{\mathbb{V}, \phi}^{(n)}$.
3. If $d_G \geq \frac{\epsilon n}{\lambda v(\mathbb{V})}$, return an arbitrary value.
4. Else, use the algorithm from [Theorem 4.2](#) to $\frac{\epsilon}{3}$ -approximate $Z_G\left(\frac{\lambda v(\mathbb{V})}{n}\right)$ with an error probability of at most $\frac{1}{9}$ and return the result.

We proceed by arguing that this procedure yields an ε -approximation of $\Xi_{\mathbb{V}}(\lambda, \phi)$ in time $\text{poly}(v(\mathbb{V})\varepsilon^{-1})$. We start by bounding the probability that the computed value is not an ε -approximation.

First, we assume that, whenever $d_G \geq \frac{en}{\lambda v(\mathbb{V})}$, the algorithm returns no ε -approximation in step 3. Let A be the event that this happens. Second, let B denote the event that the hard-core partition function $Z_G\left(\frac{\lambda v(\mathbb{V})}{n}\right)$ of the graph G that we drew in step 2 is not an $\frac{\varepsilon}{3}$ -approximation of $\Xi_{\mathbb{V}}(\lambda, \phi)$. Finally, let C denote the event we do not manage to compute an $\frac{\varepsilon}{3}$ -approximation of $Z_G\left(\frac{\lambda v(\mathbb{V})}{n}\right)$ in step 4. Note that the probability that the above procedure does not output an ε -approximation for $\Xi_{\mathbb{V}}(\lambda, \phi)$ is upper bounded by

$$\Pr\left[A \cup (B \cap \bar{A}) \cup (C \cap \bar{B} \cap \bar{A})\right] \leq \Pr[A] + \Pr[B] + \Pr[C].$$

We proceed with bounding each of these probabilities separately.

To bound $\Pr[A]$, let $z = 24 \max\left\{\frac{1}{e - \lambda C_\phi}, \frac{\lambda C_\phi}{(e - \lambda C_\phi)^2}\right\} \lambda v(\mathbb{V})$. As we are interested in asymptotic behavior in terms of $v(\mathbb{V})$, we may assume that $v(\mathbb{V})$ is sufficiently large to ensure $z \geq 5$. Note that for this, we have to exclude the case $\lambda = 0$, which trivially yields $\Xi_{\mathbb{V}}(\lambda, \phi) = 1$. Now, observe that for $z \geq 5$ it holds that $z \ln(z)^2 \geq z \ln(z \ln(z)^2)$. Next, observe that $n \geq z \ln(z)^2$. Thus, we have $n \geq z \ln(n)$. Furthermore, by $n \geq 5 \ln(5)^2 \geq e \geq 2$, we have

$$\begin{aligned} n - 1 &\geq \frac{n}{2} \\ &\geq 12 \max\left\{\frac{1}{e - \lambda C_\phi}, \frac{\lambda C_\phi}{(e - \lambda C_\phi)^2}\right\} \lambda v(\mathbb{V}) \ln(n) \\ &\geq 3(\ln(9) + 1) \max\left\{\frac{1}{e - \lambda C_\phi}, \frac{\lambda C_\phi}{(e - \lambda C_\phi)^2}\right\} \lambda v(\mathbb{V}) \ln(n) \\ &= 3(\ln(9) \ln(n) + \ln(n)) \max\left\{\frac{1}{e - \lambda C_\phi}, \frac{\lambda C_\phi}{(e - \lambda C_\phi)^2}\right\} \lambda v(\mathbb{V}) \\ &\geq 3 \ln(9n) \max\left\{\frac{1}{e - \lambda C_\phi}, \frac{\lambda C_\phi}{(e - \lambda C_\phi)^2}\right\} \lambda v(\mathbb{V}). \end{aligned}$$

Thus, we obtain

$$n \geq 3 \max\left\{\frac{\lambda C_\phi}{e - \lambda C_\phi}, \left(\frac{\lambda C_\phi}{e - \lambda C_\phi}\right)^2\right\} \ln(9n) C_\phi^{-1} v(\mathbb{V}) + 1$$

and by [Lemma 4.4](#)

$$\Pr\left[d_G \geq \frac{en}{\lambda v(\mathbb{V})}\right] \leq \Pr\left[d_G \geq \left(1 + \frac{e - \lambda C_\phi}{\lambda C_\phi}\right) \frac{n - 1}{v(\mathbb{V})} C_\phi\right] \leq \frac{1}{9}.$$

To bound $\Pr[B]$, note that for $n \geq 324e^{-2} \max\{e^6 \lambda^2 v(\mathbb{V})^2, \ln(4e^{-1})^2\}$ [Theorem 1.1](#) yields

$$\Pr[B] = \Pr\left[\left|Z_G\left(\frac{\lambda v(\mathbb{V})}{n}\right) - \Xi_{\mathbb{V}}(\lambda, \phi)\right| \geq \frac{\varepsilon}{3} \Xi_{\mathbb{V}}(\lambda, \phi)\right] \leq \frac{1}{9}.$$

Finally, note that, by [Theorem 4.2](#), we can obtain an $\frac{\varepsilon}{3}$ -approximation of $Z_G\left(\frac{\lambda v(\mathbb{V})}{n}\right)$ with error probability

at most $\Pr[C] \leq \frac{1}{9}$ in time $\tilde{O}(n^2 \varepsilon^{-2})$ as long as $\frac{\lambda \nu(\mathbb{V})}{n} < \gamma_c(d_G)$. As we only run the approximation for graphs G with $d_G < \frac{\varepsilon n}{\lambda \nu(\mathbb{V})}$ it holds that

$$\frac{\lambda \nu(\mathbb{V})}{n} < \frac{\varepsilon}{d_G} < \gamma_c(d_G),$$

proving that the requirement is satisfied.

We obtain that the error probability is bounded by $\frac{1}{3}$. To finish the proof, we need to argue that our algorithm has the desired running time. To this end, note that $N \in O(\nu(\mathbb{V})^2 \varepsilon^{-2})$. Thus, we can also choose $n \in O(\nu(\mathbb{V})^2 \varepsilon^{-2})$. By assumption, step 2 can be computed in time $t_{\mathbb{V}, \phi}(n) = t_{\mathbb{V}, \phi}(O(\nu(\mathbb{V})^2 \varepsilon^{-2}))$. Furthermore, step 3 can be computed in time $\tilde{O}(n^2 \nu(\mathbb{V})^{-1}) = \tilde{O}(\nu(\mathbb{V})^3 \varepsilon^{-4})$ and, by [Theorem 4.2](#), step 4 runs in time $\tilde{O}(n^2 \varepsilon^{-2}) = \tilde{O}(\nu(\mathbb{V})^4 \varepsilon^{-6})$ for $\frac{\lambda \nu(\mathbb{V})}{n} < \gamma_c(d_G)$. Consequently, the overall running time is in $\tilde{O}(\nu(\mathbb{V})^4 \varepsilon^{-6}) + t_{\mathbb{V}, \phi}(O(\nu(\mathbb{V})^2 \varepsilon^{-2}))$. \blacksquare

5 Sampling from repulsive Gibbs point processes

In this section, we propose an approximate sampling algorithm for the Gibbs measure of a repulsive Gibbs point process, based in random hard-core models. More precisely, we investigate the sampling procedure given by [Algorithm 1](#)

Algorithm 1: Approximate sampling algorithm for a repulsive point process $(\mathbb{V}, \lambda, \phi)$.

Data: Instance of a repulsive Gibbs point process $(\mathbb{V}, \lambda, \phi)$, error bound $\varepsilon \in (0, 1]$

Result: multiset of points in \mathbb{V}

- 1 set $n = \left\lceil \max \left\{ \begin{array}{l} 8 \frac{18^2 \cdot 12}{\varepsilon^3} \max\{e^6 \lambda^2 \nu(\mathbb{V})^2, \ln(\frac{4 \cdot 18}{\varepsilon})\}, \\ 6 \ln(\frac{4\varepsilon}{\varepsilon}) \max\left\{ \frac{1}{\varepsilon - \lambda C_\phi}, \frac{\lambda C_\phi}{(e - \lambda C_\phi)^2} \right\} \lambda \nu(\mathbb{V}) \ln \left(3 \ln(\frac{4\varepsilon}{\varepsilon}) \max\left\{ \frac{1}{\varepsilon - \lambda C_\phi}, \frac{\lambda C_\phi}{(e - \lambda C_\phi)^2} \right\} \lambda \nu(\mathbb{V}) \right)^2 \end{array} \right\rceil$;
- 2 for each $i \in [n]$ draw $X_i \sim u_{\mathbb{V}}$ independently;
- 3 draw $E \subseteq \binom{[n]}{2}$ s.t. $\{i, j\} \in E$ with probability $W_\phi(X_i, X_j) = 1 - e^{-\phi(X_i, X_j)}$ independently;
- 4 set $G = ([n], E)$;
- 5 **if** maximum degree $d_G \geq \frac{\varepsilon n}{\lambda \nu(\mathbb{V})}$ **then**
- 6 | set $X = \emptyset$;
- 7 **else**
- 8 | sample $\sigma \in \Sigma_G$ $\frac{\varepsilon}{4}$ -approximately from the hard-core distribution $\mu_G^{(\gamma(n))}$ where $\gamma(n) = \frac{\lambda \nu(\mathbb{V})}{n}$;
- 9 | set $X = \{X_i \mid i \in [n] \text{ s.t. } \sigma(i) = 1\}$ (possibly multiset);
- 10 **end**
- 11 **return** X ;

Our main theorem in this section is as follows.

► **Theorem 5.1.** Let (\mathbb{X}, d) be a complete, separable metric space, let $\mathcal{B} = \mathcal{B}(\mathbb{X})$ be the Borel algebra and let ν be a locally finite reference measure on $(\mathbb{X}, \mathcal{B})$. Let $\mathbb{V} \subseteq \mathbb{X}$ be bounded and measurable, let $\lambda \in \mathbb{R}_{\geq 0}$ and let $\phi : \mathbb{X}^2 \rightarrow \mathbb{R}_{\geq 0} \cup \{\infty\}$ be a symmetric repulsive potential. Assume we can sample from the uniform distribution $u_{\mathbb{V}}$ in time $t_{\mathbb{V}}$ and, for every $x, y \in \mathbb{V}$, evaluate $\phi(x, y)$ in time t_ϕ . If the Gibbs point process $P_{\mathbb{V}}^{(\lambda, \phi)}$ is simple and $\lambda < \frac{\varepsilon}{C_\phi}$ then, for every $\varepsilon \in \mathbb{R}_{> 0}$, [Algorithm 1](#) samples ε -approximately from $P_{\mathbb{V}}^{(\lambda, \phi)}$ and has running time in $\tilde{O}(\nu(\mathbb{V})^2 \varepsilon^{-4} + \nu(\mathbb{V})^2 \varepsilon^{-3} t_{\mathbb{V}} + \nu(\mathbb{V})^4 \varepsilon^{-6} t_\phi)$. \blacktriangleleft

[Theorem 1.3](#) follows immediately from the theorem above. To prove [Theorem 5.1](#), we start by analyzing a simplified algorithm, given in [Algorithm 2](#).

Algorithm 2: Modified sampling process

- Data:** Instance of a repulsive Gibbs point process $(\mathbb{V}, \lambda, \phi)$, error bound $\varepsilon \in (0, 1]$
- Result:** multiset of points in \mathbb{V}
- 1 set $n = \left\lceil \max \left\{ \frac{8 \frac{18^2 - 12}{\varepsilon^3} \max\{e^{6\lambda^2 \nu(\mathbb{V})^2}, \ln(\frac{4 \cdot 18}{\varepsilon})\}, 6 \ln(\frac{4\varepsilon}{\varepsilon}) \max\left\{\frac{1}{e^{-\lambda C_\phi}}, \frac{\lambda C_\phi}{(e^{-\lambda C_\phi})^2}\right\} \lambda \nu(\mathbb{V}) \ln\left(3 \ln(\frac{4\varepsilon}{\varepsilon}) \max\left\{\frac{1}{e^{-\lambda C_\phi}}, \frac{\lambda C_\phi}{(e^{-\lambda C_\phi})^2}\right\} \lambda \nu(\mathbb{V})\right)^2 \right\} \right\rceil$;
 - 2 for each $i \in [n]$ draw $X_i \sim u_{\mathbb{V}}$ independently;
 - 3 draw $E \subseteq \binom{[n]}{2}$ s.t. $\{i, j\} \in E$ with probability $W_\phi(X_i, X_j) = 1 - e^{-\phi(X_i, X_j)}$ independently;
 - 4 set $G = ([n], E)$;
 - 5 sample $\tau \in \Sigma_G$ exactly from the hard-core distribution $\mu_G^{(\gamma(n))}$ where $\gamma(n) = \frac{\lambda \nu(\mathbb{V})}{n}$;
 - 6 set $Y = \{X_i \mid i \in [n] \text{ s.t. } \tau(i) = 1\}$ (possibly multiset);
 - 7 **return** Y ;
-

The main difference between [Algorithm 1](#) and [Algorithm 2](#) is that the latter one does not check if the maximum degree of the sampled graph G is bounded and that it assumes access to a perfect sampler for $\mu_G^{(\gamma(n))}$. It is not clear if such a perfect sampler for the hard-core Gibbs distribution can be realized in polynomial time, especially for arbitrary vertex degrees. Therefore, [Algorithm 2](#) is not suitable for algorithmic applications. However, the main purpose of [Algorithm 2](#) is that the distribution of point multisets that it outputs are much easier to analyze. We use this, together with a coupling argument, to bound the total variation distance between the output of [Algorithm 1](#) and $P_{\mathbb{V}}^{(\lambda, \phi)}$. Once this is done, it remains to show that [Algorithm 1](#) satisfies the running time requirements, given in [Theorem 5.1](#).

To analyze the output distribution of [Algorithm 2](#), we start by considering the resulting distribution of multisets of points (or counting measures respectively) when conditioning on the event that the hard-core partition function $Z_G(\gamma(n))$ of the drawn graph G is close to the partition function of the continuous process $\Xi_{\mathbb{V}}(\lambda, \phi)$. More specifically, for any given n and $\alpha \in \mathbb{R}_{\geq 0}$, let $A_\alpha^{(n)} = \{H \in \mathcal{G}_n \mid |Z_H(\gamma(n)) - \Xi_{\mathbb{V}}(\lambda, \phi)| \leq \alpha \Xi_{\mathbb{V}}(\lambda, \phi)\}$. We derive an explicit density for the output of [Algorithm 2](#) with respect to a Poisson point process under the condition that $G \in A_\alpha^{(n)}$ for some sufficiently small α . To this end, we use the following characterization of simple point processes via so called *void probabilities*.

► **Theorem 5.2 (Rényi–Mönch, see [13, Theorem 9.2.XII]).** Let (\mathbb{X}, d) be a complete, separable metric space, let $\mathcal{B} = \mathcal{B}(\mathbb{X})$ be the associated Borel algebra. Let P and Q be simple point process on (\mathbb{X}, d) . If, for $\eta_P \sim P$ and $\eta_Q \sim Q$ and for all bounded $B \in \mathcal{B}$, it holds that

$$\Pr[\eta_P(B) = 0] = \Pr[\eta_Q(B) = 0],$$

then $P = Q$. ◀

[Theorem 5.2](#) greatly simplifies proving that a given candidate function actually is a valid density for the point process in question, as it implies that it is sufficient to check if it yields the correct void probabilities.

Before we proceed, we introduce some additional notation that is useful for stating and proving our next lemmas. For a given graph $H = (V, E)$, we denote by $\mathcal{I}(H) \subseteq 2^V$ the set of all independent sets in H . Moreover, for every spin configuration $\sigma \in \Sigma_H$, we denote by S_σ the set of all vertices $v \in V$ with $\sigma(v) = 1$. Note that, for a hard-core model on H with $\gamma > 0$, this construction gives a one-to-one correspondence

between $\mathcal{I}(H)$ and the set of spin configurations $\sigma \in \Sigma_H$ with $\mu_H^{(\gamma)}(\sigma) > 0$. Therefore, it is often convenient to argue about elements in $\mathcal{I}(H)$ instead of using spin configurations.

► **Lemma 5.3.** Let (\mathbb{X}, d) be a complete, separable metric space, let $\mathcal{B} = \mathcal{B}(\mathbb{X})$ be the Borel algebra and let ν be a locally finite reference measure on $(\mathbb{X}, \mathcal{B})$. Let $\mathbb{V} \subseteq \mathbb{X}$ be bounded and measurable, let $\lambda \in \mathbb{R}_{\geq 0}$ and let $\phi : \mathbb{X}^2 \rightarrow \mathbb{R}_{\geq 0} \cup \{\infty\}$ be a symmetric repulsive potential. Furthermore, for any given $\varepsilon \in (0, 1]$, let \widehat{P}_ε be the point process produced by [Algorithm 2](#) conditioned on $G \in A_{\frac{\varepsilon}{12}}^{(n)}$, and let Q_λ denote a Poisson point process with intensity λ . If the Gibbs point process $P_{\mathbb{V}}^{(\lambda, \phi)}$ is simple, then \widehat{P}_ε has a density with respect to Q of the form

$$g_\varepsilon(\eta) = \mathbb{1}_{\eta \in \mathcal{N}_{\mathbb{V}}} \Pr \left[G \in A_{\frac{\varepsilon}{12}}^{(n)} \right]^{-1} \left(\prod_{i=0}^{\eta(\mathbb{V})-1} 1 - \frac{i}{n} \right) \mathbb{1}_{\eta(\mathbb{V}) \leq n} \\ \cdot \left(\prod_{\{x, y\} \in \binom{X_\eta}{2}} e^{-N_x(\eta) N_y(\eta) \phi(x, y)} \right) \left(\prod_{x \in X_\eta} e^{-\frac{N_x(\eta)(N_x(\eta)-1)}{2} \phi(x, x)} \right) \Psi_n^{(\eta(\mathbb{V}))}(\varphi(\eta)) e^{\lambda \nu(\mathbb{V})},$$

where φ maps every finite counting measure η to an arbitrary but fixed tuple $(x_1, \dots, x_{\eta(\mathbb{X})})$ such that $\eta = \sum_{i=1}^{\eta(\mathbb{X})} \delta_{x_i}$ and

$$\Psi_n^{(k)}(\mathbf{x}) = \sum_{\substack{H \in A_{\frac{\varepsilon}{12}}^{(n)} \\ [k] \in \mathcal{I}(H)}} \frac{1}{Z_H(\gamma(n))} \int_{\mathbb{V}^{n-k}} \left(\prod_{\substack{(i, j) \in [k] \times [n-k] \\ \{i, j+k\} \in E_H}} 1 - e^{-\phi(x_i, y_j)} \right) \left(\prod_{\substack{(i, j) \in [k] \times [n-k] \\ \{i, j+k\} \notin E_H}} e^{-\phi(x_i, y_j)} \right) \\ \left(\prod_{\substack{\{i, j\} \in \binom{[n-k]}{2} \\ \{i+k, j+k\} \in E_H}} 1 - e^{-\phi(y_i, y_j)} \right) \left(\prod_{\substack{\{i, j\} \in \binom{[n-k]}{2} \\ \{i+k, j+k\} \notin E_H}} e^{-\phi(y_i, y_j)} \right) u_{\mathbb{V}}^{n-k}(\mathbf{d}\mathbf{y})$$

for all $\mathbf{x} = (x_1, \dots, x_k) \in \mathbb{V}^k$. ◀

Proof. First, observe that $G \sim \zeta_{\mathbb{V}, \phi}^{(n)}$. As $n \geq 4 \frac{12^3}{\varepsilon^3} \max \left\{ e^6 \lambda^2 \nu(\mathbb{V})^2, \ln \left(4 \frac{12}{\varepsilon} \right)^2 \right\}$, [Theorem 1.1](#) implies that $\Pr \left[G \in A_{\frac{\varepsilon}{12}}^{(n)} \right] \geq 1 - \frac{\varepsilon}{12} > 0$. Therefore, conditioning on the event $G \in A_{\frac{\varepsilon}{12}}^{(n)}$ is well defined.

Next, note that, for all $x \in \mathbb{V}$ it holds that

$$\Pr[\eta(\{x\}) \geq 2] \geq \frac{e^{-\phi(x, x)} \lambda^2 \nu(\{x\})^2}{\Xi_{\mathbb{V}}(\lambda, \phi)} \geq \frac{e^{-\phi(x, x)} \lambda^2 \nu(\{x\})^2}{e^{\lambda \nu(\mathbb{V})}}$$

for $\eta \sim P_{\mathbb{V}}^{(\lambda, \phi)}$. Thus, if $P_{\mathbb{V}}^{(\lambda, \phi)}$ is simple (i.e., $\Pr[\eta(\{x\}) \geq 2] = 0$ for all $x \in \mathbb{V}$), it holds that $\lambda = 0$ or, for all $x \in \mathbb{V}$, $\nu(\{x\}) = 0$ or $\phi(x, x) = \infty$. This implies that the output of [Algorithm 2](#) is simple as well, and consequently \widehat{P}_ε is a simple point process.

Knowing that \widehat{P}_ε is simple, [Theorem 5.2](#) implies that, in order to verify that g_ε is indeed a density for \widehat{P}_ε , it suffices to prove that it yields the correct void probabilities. Formally, this means showing that for all

bounded $B \in \mathcal{B}$ it holds that

$$\Pr \left[Y \cap B = \emptyset \mid G \in A_{\frac{\varepsilon}{12}}^{(n)} \right] = \int_{\mathcal{N}} \mathbb{1}_{\eta(B)=0} g_{\varepsilon}(\eta) Q_{\lambda}(\mathrm{d}\eta)$$

for Y and G as in [Algorithm 2](#).

To prove this, we first write

$$\begin{aligned} \int_{\mathcal{N}} \mathbb{1}_{\eta(B)=0} g_{\varepsilon}(\eta) Q_{\lambda}(\mathrm{d}\eta) &= \int_{\mathcal{N}_{\mathbb{V}}} \mathbb{1}_{\eta(B)=0} g_{\varepsilon}(\eta) Q_{\lambda}(\mathrm{d}\eta) \\ &= e^{-\lambda v(\mathbb{V})} \cdot \left(g_{\varepsilon}(\mathbf{0}) + \sum_{k \in \mathbb{N}_{\geq 1}} \frac{\lambda^k}{k!} \int_{\mathbb{V}^k} \mathbb{1}_{\forall i \in [k]: x_i \notin B} g_{\varepsilon} \left(\sum_{i \in [k]} \delta_{x_i} \right) v^k(\mathrm{d}\mathbf{x}) \right), \end{aligned}$$

where $\mathbf{0}$ denotes the constant 0 measure on \mathbb{X} . Note that

$$\begin{aligned} &e^{-\lambda v(\mathbb{V})} g_{\varepsilon}(\mathbf{0}) \\ &= \Pr \left[G \in A_{\frac{\varepsilon}{12}}^{(n)} \right]^{-1} \cdot \sum_{H \in A_{\frac{\varepsilon}{12}}^{(n)}} \frac{1}{Z_H(Y(n))} \int_{\mathbb{V}^n} \left(\prod_{\substack{\{i,j\} \in \binom{[n]}{2}: \\ \{i,j\} \in E_H}} 1 - e^{-\phi(y_i, y_j)} \right) \left(\prod_{\substack{\{i,j\} \in \binom{[n]}{2}: \\ \{i,j\} \notin E_H}} e^{-\phi(y_i, y_j)} \right) u_{\mathbb{V}}^n(\mathrm{d}\mathbf{y}) \\ &= \Pr \left[G \in A_{\frac{\varepsilon}{12}}^{(n)} \right]^{-1} \cdot \sum_{H \in A_{\frac{\varepsilon}{12}}^{(n)}} \Pr[S_{\tau} = \emptyset \mid G = H] \Pr[G = H] \\ &= \frac{\Pr \left[S_{\tau} = \emptyset \wedge G \in A_{\frac{\varepsilon}{12}}^{(n)} \right]}{\Pr \left[G \in A_{\frac{\varepsilon}{12}}^{(n)} \right]} \\ &= \Pr \left[S_{\tau} = \emptyset \mid G \in A_{\frac{\varepsilon}{12}}^{(n)} \right] \end{aligned}$$

for τ as in [Algorithm 2](#). We proceed by a case distinction based on k . For every $k > n$ and $(x_1, \dots, x_k) \in \mathbb{V}^k$ we have $g_{\varepsilon}(\sum_{i \in [k]} \delta_{x_k}) = 0$. Therefore, we get

$$\int_{\mathbb{V}^k} \mathbb{1}_{\forall i \in [k]: x_i \notin B} g_{\varepsilon} \left(\sum_{i \in [k]} \delta_{x_i} \right) v^k(\mathrm{d}\mathbf{x}) = 0$$

for all $k > n$. Now, consider $k \in [n]$ and observe that for all $\mathbf{x} = (x_1, \dots, x_k) \in \mathbb{V}^k$ we have

$$\Psi_n^{(k)} \left(\varphi \left(\sum_{i \in [k]} \delta_{x_i} \right) \right) = \Psi_n^{(k)}(\mathbf{x})$$

by symmetry. Moreover, it holds that

$$\frac{\lambda^k}{k!} \left(\prod_{i=0}^{k-1} 1 - \frac{i}{n} \right) = \binom{n}{k} \frac{\gamma(n)^k}{\nu(\mathbb{V})^k}.$$

Therefore, we have

$$\begin{aligned} & e^{-\lambda\nu(\mathbb{V})} \frac{\lambda^k}{k!} \int_{\mathbb{V}^k} \mathbb{1}_{\forall i \in [k]: x_i \notin B} \mathcal{G}_\varepsilon \left(\sum_{i \in [k]} \delta_{x_i} \right) \nu^k(\mathbf{d}\mathbf{x}) \\ &= \Pr \left[G \in A_{\frac{\varepsilon}{12}}^{(n)} \right]^{-1} \binom{n}{k} \gamma(n)^k \int_{\mathbb{V}^k} \mathbb{1}_{\forall i \in [k]: x_i \notin B} \left(\prod_{i,j \in \binom{[k]}{2}} e^{-\phi(x_i, x_j)} \right) \Psi_n^{(k)}(\mathbf{x}) u_{\mathbb{V}}^k(\mathbf{d}\mathbf{x}). \end{aligned}$$

Next, note that

$$\begin{aligned} & \gamma(n)^k \int_{\mathbb{V}^k} \mathbb{1}_{\forall i \in [k]: x_i \notin B} \left(\prod_{i,j \in \binom{[k]}{2}} e^{-\phi(x_i, x_j)} \right) \Psi_n^{(k)}(\mathbf{x}) u_{\mathbb{V}}^k(\mathbf{d}\mathbf{x}) \\ &= \sum_{H \in A_{\frac{\varepsilon}{12}}^{(n)}} \mathbb{1}_{[k] \in I(H)} \frac{\gamma(n)^k}{Z_H(\gamma(n))} \int_{\mathbb{V}^n} \mathbb{1}_{\forall i \in [k]: x_i \notin B} \left(\prod_{\substack{\{i,j\} \in \binom{[n]}{2}: \\ \{i,j\} \in E_H}} 1 - e^{-\phi(x_i, x_j)} \right) \left(\prod_{\substack{\{i,j\} \in \binom{[n]}{2}: \\ \{i,j\} \notin E_H}} e^{-\phi(x_i, x_j)} \right) u_{\mathbb{V}}^n(\mathbf{d}\mathbf{x}) \\ &= \sum_{H \in A_{\frac{\varepsilon}{12}}^{(n)}} \Pr[S_\tau = [k] \mid G = H] \Pr[G = H \wedge \forall i \in [k] : X_i \notin B] \end{aligned}$$

for X_1, \dots, X_n as in [Algorithm 2](#). Furthermore, because the event $S_\tau = [k]$ is independent of X_1, \dots, X_n given G , it holds that

$$\begin{aligned} & \sum_{H \in A_{\frac{\varepsilon}{12}}^{(n)}} \Pr[S_\tau = [k] \mid G = H] \Pr[G = H \wedge \forall i \in [k] : X_i \notin B] \\ &= \sum_{H \in A_{\frac{\varepsilon}{12}}^{(n)}} \Pr[S_\tau = [k] \wedge G = H \wedge \forall i \in [k] : X_i \notin B] \\ &= \Pr \left[S_\tau = [k] \wedge G \in A_{\frac{\varepsilon}{12}}^{(n)} \wedge \forall i \in [k] : X_i \notin B \right] \end{aligned}$$

and

$$\begin{aligned} e^{-\lambda\nu(\mathbb{V})} \frac{\lambda^k}{k!} \int_{\mathbb{V}^k} \mathbb{1}_{\forall i \in [k]: x_i \notin B} \mathcal{G}_\varepsilon \left(\sum_{i \in [k]} \delta_{x_i} \right) \nu^k(\mathbf{d}\mathbf{x}) &= \binom{n}{k} \frac{\Pr \left[S_\tau = [k] \wedge G \in A_{\frac{\varepsilon}{12}}^{(n)} \wedge \forall i \in [k] : X_i \notin B \right]}{\Pr \left[G \in A_{\frac{\varepsilon}{12}}^{(n)} \right]} \\ &= \binom{n}{k} \Pr \left[S_\tau = [k] \wedge \forall i \in [k] : X_i \notin B \mid G \in A_{\frac{\varepsilon}{12}}^{(n)} \right] \end{aligned}$$

$$= \sum_{V' \in \binom{[n]}{k}} \Pr \left[S_\tau = V' \wedge \forall i \in V' : X_i \notin B \mid G \in A_{\frac{\varepsilon}{12}}^{(n)} \right],$$

where the last equality is due to symmetry. Combining everything yields

$$\begin{aligned} \int_{\mathcal{N}} \mathbb{1}_{\eta(B)=0} g_\varepsilon(\eta) Q_\lambda(d\eta) &= \Pr \left[S_\tau = \emptyset \mid G \in A_{\frac{\varepsilon}{12}}^{(n)} \right] + \sum_{k=1}^n \sum_{V' \in \binom{[n]}{k}} \Pr \left[S_\tau = V' \wedge \forall i \in V' : X_i \notin B \mid G \in A_{\frac{\varepsilon}{12}}^{(n)} \right] \\ &= \sum_{V' \in 2^{[n]}} \Pr \left[S_\tau = V' \wedge \forall i \in V' : X_i \notin B \mid G \in A_{\frac{\varepsilon}{12}}^{(n)} \right] \\ &= \Pr \left[\forall i \in S_\tau : X_i \notin B \mid G \in A_{\frac{\varepsilon}{12}}^{(n)} \right] \\ &= \Pr \left[Y \cap B = \emptyset \mid G \in A_{\frac{\varepsilon}{12}}^{(n)} \right], \end{aligned}$$

which concludes the proof. ■

We proceed by upper and lower bounding the density $g_\varepsilon(\eta)$ in terms of the density of $P_{\mathbb{V}}^{(\lambda, \phi)}$. To this end, we use the following basic facts about the partition function of the hard-core model.

► **Observation 5.4** (see [17]). For every undirected graph $G = (V, E)$ the following holds:

1. For all $\gamma_1, \gamma_2 \in \mathbb{R}_{\geq 0}$

$$Z_G(\gamma_1) \leq Z_G(\gamma_1 + \gamma_2) \leq e^{\gamma_2 |V|} Z_G(\gamma_1).$$

2. For all $\gamma \in \mathbb{R}_{\geq 0}$ and $S \subseteq V$

$$Z_{G-S}(\gamma) \leq Z_G(\gamma) \leq e^{\gamma |S|} Z_{G-S}(\gamma),$$

where $G - S$ denotes the subgraph of G that is induced by $V \setminus S$. ◀

Using [Observation 5.4](#) we derive the following bounds.

► **Lemma 5.5.** Consider the setting of [Lemma 5.3](#) and let f denote the density of $P_{\mathbb{V}}^{(\lambda, \phi)}$ with respect to Q_λ . For n as in [Algorithm 2](#) and all $\eta \in \mathcal{N}$ with $\eta(\mathbb{V}) \leq \min \left\{ \sqrt{\frac{\varepsilon n}{12}}, \frac{\varepsilon}{40\lambda \nu(\mathbb{V}) + \varepsilon} n \right\}$ it holds that

$$\left(1 - \frac{\varepsilon}{4}\right) f(\eta) \leq g_\varepsilon(\eta) \leq \left(1 + \frac{\varepsilon}{4}\right) f(\eta). \quad \text{◀}$$

Proof. First, recall that, when $P_{\mathbb{V}}^{(\lambda, \phi)}$ is simple, its density with respect to Q_λ can be expressed as

$$f(\eta) = \frac{1}{\Xi_{\mathbb{V}}(\lambda, \phi)} \mathbb{1}_{\eta \in \mathcal{N}_{\mathbb{V}}} \left(\prod_{\{x, y\} \in \binom{X_\eta}{2}} e^{-N_x(\eta) N_y(\eta) \phi(x, y)} \right) \left(\prod_{x \in X_\eta} e^{-\frac{N_x(\eta)(N_x(\eta)-1)}{2} \phi(x, x)} \right) e^{\lambda \nu(\mathbb{V})}$$

for every $\eta \in \mathcal{N}$. Therefore, we have

$$g_\varepsilon(\eta) = \Pr \left[G \in A_{\frac{\varepsilon}{12}}^{(n)} \right]^{-1} \left(\prod_{i=0}^{\eta(\mathbb{V})-1} 1 - \frac{i}{n} \right) \mathbb{1}_{\eta(\mathbb{V}) \leq n} \Psi_n^{(\eta(\mathbb{V}))}(\varphi(\eta)) \Xi_{\mathbb{V}}(\lambda, \phi) f(\eta).$$

As we focus on η with $\eta(\mathbb{V}) \leq \sqrt{\frac{\varepsilon n}{12}} \leq n$, we omit the indicator $\mathbb{1}_{\eta(\mathbb{V}) \leq n}$ from now on.

We proceed by deriving an upper bound on $g_\varepsilon(\eta)$ for $\eta \in \mathcal{N}$ with $\eta(\mathbb{V}) \leq \min\left\{\sqrt{\frac{\varepsilon n}{12}}, \frac{\varepsilon}{40\lambda v(\mathbb{V}) + \varepsilon}n\right\}$. To this end, note that

$$\left(\prod_{i=0}^{\eta(\mathbb{V})-1} 1 - \frac{i}{n}\right) \leq 1.$$

Moreover, for $G \sim \zeta_{\mathbb{V}, \phi}^{(n)}$ and $n \geq 4\frac{12^3}{\varepsilon^3} \max\left\{e^6 \lambda^2 v(\mathbb{V})^2, \ln\left(4\frac{12}{\varepsilon}\right)^2\right\}$ [Theorem 1.1](#) yields $\Pr\left[G \in A_{\frac{\varepsilon}{12}}^{(n)}\right] \geq 1 - \frac{\varepsilon}{12}$. Finally, observe that for all $\mathbf{x} \in \mathbb{V}^k$ for $k \leq n$ we have

$$\begin{aligned} \Psi_n^{(k)}(\mathbf{x}) &\leq \frac{1}{\left(1 - \frac{\varepsilon}{12}\right)\Xi_{\mathbb{V}}(\lambda, \phi)} \sum_{\substack{H \in A_{\frac{\varepsilon}{12}}^{(n)}: \\ [k] \in \mathcal{I}(H)}} \int_{\mathbb{V}^{n-k}} \left(\prod_{\substack{(i,j) \in [k] \times [n-k]: \\ \{i,j+k\} \in E_H}} 1 - e^{-\phi(x_i, y_j)} \right) \left(\prod_{\substack{(i,j) \in [k] \times [n-k]: \\ \{i,j+k\} \notin E_H}} e^{-\phi(x_i, y_j)} \right) \\ &\quad \cdot \left(\prod_{\substack{\{i,j\} \in \binom{[n-k]}{2}: \\ \{i+k, j+k\} \in E_H}} 1 - e^{-\phi(y_i, y_j)} \right) \left(\prod_{\substack{\{i,j\} \in \binom{[n-k]}{2}: \\ \{i+k, j+k\} \notin E_H}} e^{-\phi(y_i, y_j)} \right) u_{\mathbb{V}}^{n-k}(\mathbf{d}\mathbf{y}) \\ &= \frac{1}{\left(1 - \frac{\varepsilon}{12}\right)\Xi_{\mathbb{V}}(\lambda, \phi)} \int_{\mathbb{V}^{n-k}} \sum_{\substack{H \in A_{\frac{\varepsilon}{12}}^{(n)}: \\ [k] \in \mathcal{I}(H)}} \left(\prod_{\substack{(i,j) \in [k] \times [n-k]: \\ \{i,j+k\} \in E_H}} 1 - e^{-\phi(x_i, y_j)} \right) \left(\prod_{\substack{(i,j) \in [k] \times [n-k]: \\ \{i,j+k\} \notin E_H}} e^{-\phi(x_i, y_j)} \right) \\ &\quad \cdot \left(\prod_{\substack{\{i,j\} \in \binom{[n-k]}{2}: \\ \{i+k, j+k\} \in E_H}} 1 - e^{-\phi(y_i, y_j)} \right) \left(\prod_{\substack{\{i,j\} \in \binom{[n-k]}{2}: \\ \{i+k, j+k\} \notin E_H}} e^{-\phi(y_i, y_j)} \right) u_{\mathbb{V}}^{n-k}(\mathbf{d}\mathbf{y}) \\ &\leq \frac{1}{\left(1 - \frac{\varepsilon}{12}\right)\Xi_{\mathbb{V}}(\lambda, \phi)} \int_{\mathbb{V}^{n-k}} 1 u_{\mathbb{V}}^{n-k}(\mathbf{d}\mathbf{y}) \\ &\leq \frac{1}{\left(1 - \frac{\varepsilon}{12}\right)\Xi_{\mathbb{V}}(\lambda, \phi)}. \end{aligned}$$

Given that $\varepsilon \leq 1$ we get

$$g_\varepsilon(\eta) \leq \left(1 - \frac{\varepsilon}{12}\right)^{-2} f(\eta) \leq \left(1 + \frac{\varepsilon}{11}\right)^2 f(\eta) \leq \left(1 + \frac{\varepsilon}{4}\right) f(\eta),$$

which proves the upper bound.

For the lower bound, note that

$$\Pr\left[G \in A_{\frac{\varepsilon}{12}}^{(n)}\right]^{-1} \geq 1$$

and for $\eta(\mathbb{V}) \leq \sqrt{\frac{\varepsilon n}{12}}$

$$\left(\prod_{i=0}^{\eta(\mathbb{V})-1} 1 - \frac{i}{n} \right) \geq \left(1 - \frac{\eta(\mathbb{V})}{n} \right)^{\eta(\mathbb{V})} \geq 1 - \frac{\eta(\mathbb{V})^2}{n} \geq 1 - \frac{\varepsilon}{12}.$$

We proceed by lower bounding $\Psi_n^{(k)}(\mathbf{x})$. First, observe that

$$\begin{aligned} \Psi_n^{(k)}(\mathbf{x}) &\geq \frac{1}{\left(1 + \frac{\varepsilon}{12}\right) \Xi_{\mathbb{V}}(\lambda, \phi)} \sum_{\substack{H \in \mathcal{A}_{\frac{\varepsilon}{12}}^{(n)}: \\ [k] \in \mathcal{I}(H)}} \int_{\mathbb{V}^{n-k}} \left(\prod_{\substack{(i,j) \in [k] \times [n-k]: \\ \{i,j+k\} \in E_H}} 1 - e^{-\phi(x_i, y_j)} \right) \left(\prod_{\substack{(i,j) \in [k] \times [n-k]: \\ \{i,j+k\} \notin E_H}} e^{-\phi(x_i, y_j)} \right) \\ &\quad \cdot \left(\prod_{\substack{\{i,j\} \in \binom{[n-k]}{2}: \\ \{i+k, j+k\} \in E_H}} 1 - e^{-\phi(y_i, y_j)} \right) \left(\prod_{\substack{\{i,j\} \in \binom{[n-k]}{2}: \\ \{i+k, j+k\} \notin E_H}} e^{-\phi(y_i, y_j)} \right) u_{\mathbb{V}}^{n-k}(\mathbf{d}\mathbf{y}) \\ &= \frac{1}{\left(1 + \frac{\varepsilon}{12}\right) \Xi_{\mathbb{V}}(\lambda, \phi)} \sum_{\substack{H \in \mathcal{G}_n: \\ [k] \in \mathcal{I}(H)}} \mathbb{1}_{H \in \mathcal{A}_{\frac{\varepsilon}{12}}^{(n)}} \int_{\mathbb{V}^{n-k}} \left(\prod_{\substack{(i,j) \in [k] \times [n-k]: \\ \{i,j+k\} \in E_H}} 1 - e^{-\phi(x_i, y_j)} \right) \left(\prod_{\substack{(i,j) \in [k] \times [n-k]: \\ \{i,j+k\} \notin E_H}} e^{-\phi(x_i, y_j)} \right) \\ &\quad \cdot \left(\prod_{\substack{\{i,j\} \in \binom{[n-k]}{2}: \\ \{i+k, j+k\} \in E_H}} 1 - e^{-\phi(y_i, y_j)} \right) \left(\prod_{\substack{\{i,j\} \in \binom{[n-k]}{2}: \\ \{i+k, j+k\} \notin E_H}} e^{-\phi(y_i, y_j)} \right) u_{\mathbb{V}}^{n-k}(\mathbf{d}\mathbf{y}). \end{aligned}$$

Next, for each graph $H \in \mathcal{G}_n$, let $H' = ([n-k], E')$ denote the subgraph that results from $H - [k]$ after relabeling each vertex in $i \in [n] \setminus [k]$ to $i-k \in [n-k]$ (note that this relabeling is formally required for $H' \in \mathcal{G}_{n-k}$). By [Observation 5.4](#) and the fact that $\gamma(n) \leq \gamma(n-k)$ and $Z_H(\gamma) = Z_{H-[k]}(\gamma)$ for all $\gamma \in \mathbb{R}_{\geq 0}$ we have

$$Z_H(\gamma(n)) \leq e^{\gamma(n)k} Z_{H'}(\gamma(n)) \leq e^{\frac{k}{n} \lambda_V(\mathbb{V})} Z_{H'}(\gamma(n-k)).$$

On the other hand, note that

$$\begin{aligned} \gamma(n-k) &= \frac{\lambda_V(\mathbb{V})}{n-k} = \frac{n}{n(n-k)} \lambda_V(\mathbb{V}) = \left(\frac{n-k}{n(n-k)} + \frac{k}{n(n-k)} \right) \lambda_V(\mathbb{V}) \\ &= \left(\frac{1}{n} + \frac{k}{n(n-k)} \right) \lambda_V(\mathbb{V}) = \gamma(n) + \frac{k}{n(n-k)} \lambda_V(\mathbb{V}). \end{aligned}$$

Therefore, [Observation 5.4](#) yields

$$Z_H(\gamma(n-k)) \leq e^{\frac{k}{n-k} \lambda_V(\mathbb{V})} Z_H(\gamma(n))$$

and

$$Z_H(\gamma(n)) \geq e^{-\frac{k}{n-k} \lambda_V(\mathbb{V})} Z_H(\gamma(n-k)) \geq e^{-\frac{k}{n-k} \lambda_V(\mathbb{V})} Z_{H'}(\gamma(n-k)).$$

Thus, for $k \leq \frac{\varepsilon}{40\lambda v(\mathbb{V})+\varepsilon}n$ we have

$$e^{-\frac{\varepsilon}{40}}Z_{H'}(\gamma(n-k)) \leq Z_H(\gamma(n)) \leq e^{\frac{\varepsilon}{40}}Z_{H'}(\gamma(n-k)).$$

As $e^{-\frac{\varepsilon}{40}}(1 - \frac{\varepsilon}{18}) \geq (1 - \frac{\varepsilon}{12})$ and $e^{\frac{\varepsilon}{40}}(1 + \frac{\varepsilon}{18}) \leq (1 + \frac{\varepsilon}{12})$ for all $\varepsilon \in [0, 1]$, this means that $H' \in A_{\frac{\varepsilon}{18}}^{(n-k)}$ is a sufficient condition for $H \in A_{\frac{\varepsilon}{12}}^{(n)}$ and

$$\begin{aligned} & \sum_{\substack{H \in \mathcal{G}_n: \\ [k] \in \mathcal{I}(H)}} \mathbb{1}_{H \in A_{\frac{\varepsilon}{12}}^{(n)}} \int_{\mathbb{V}^{n-k}} \left(\prod_{\substack{(i,j) \in [k] \times [n-k]: \\ \{i,j+k\} \in E_H}} 1 - e^{-\phi(x_i, y_j)} \right) \left(\prod_{\substack{(i,j) \in [k] \times [n-k]: \\ \{i,j+k\} \notin E_H}} e^{-\phi(x_i, y_j)} \right) \\ & \quad \left(\prod_{\substack{\{i,j\} \in \binom{[n-k]}{2}: \\ \{i+k, j+k\} \in E_H}} 1 - e^{-\phi(y_i, y_j)} \right) \left(\prod_{\substack{\{i,j\} \in \binom{[n-k]}{2}: \\ \{i+k, j+k\} \notin E_H}} e^{-\phi(y_i, y_j)} \right) u_{\mathbb{V}}^{n-k}(\mathbf{d}\mathbf{y}) \\ & \geq \sum_{H' \in \mathcal{G}_{n-k}} \mathbb{1}_{H' \in A_{\frac{\varepsilon}{18}}^{(n-k)}} \int_{\mathbb{V}^{n-k}} \left(\prod_{\substack{\{i,j\} \in \binom{[n-k]}{2}: \\ \{i,j\} \in E_{H'}}} 1 - e^{-\phi(y_i, y_j)} \right) \left(\prod_{\substack{\{i,j\} \in \binom{[n-k]}{2}: \\ \{i,j\} \notin E_{H'}}} e^{-\phi(y_i, y_j)} \right) \\ & \quad \sum_{F \subseteq [k] \times [n-k]} \left(\prod_{\substack{(i,j) \in [k] \times [n-k]: \\ (i,j) \in F}} 1 - e^{-\phi(x_i, y_j)} \right) \left(\prod_{\substack{(i,j) \in [k] \times [n-k]: \\ (i,j) \notin F}} e^{-\phi(x_i, y_j)} \right) u_{\mathbb{V}}^{n-k}(\mathbf{d}\mathbf{y}) \\ & = \sum_{H' \in \mathcal{G}_{n-k}} \mathbb{1}_{H' \in A_{\frac{\varepsilon}{18}}^{(n-k)}} \int_{\mathbb{V}^{n-k}} \left(\prod_{\substack{\{i,j\} \in \binom{[n-k]}{2}: \\ \{i,j\} \in E_{H'}}} 1 - e^{-\phi(y_i, y_j)} \right) \left(\prod_{\substack{\{i,j\} \in \binom{[n-k]}{2}: \\ \{i,j\} \notin E_{H'}}} e^{-\phi(y_i, y_j)} \right) u_{\mathbb{V}}^{n-k}(\mathbf{d}\mathbf{y}) \\ & = \Pr \left[G' \in A_{\frac{\varepsilon}{18}}^{(n-k)} \right] \end{aligned}$$

for $G' \sim \zeta_{\mathbb{V}, \phi}^{(n-k)}$. Next, observe that $n \geq 1$ we have $k \leq \min \left\{ \sqrt{\frac{\varepsilon n}{12}}, \frac{\varepsilon}{40\lambda v(\mathbb{V})+\varepsilon}n \right\} \leq \frac{n}{2}$ and $n-k \geq \frac{n}{2}$. Therefore, for $n \geq 8 \frac{18^2 \cdot 12}{\varepsilon^3} \max \left\{ e^6 \lambda^2 v(\mathbb{V})^2, \ln \left(\frac{4 \cdot 18}{\varepsilon} \right)^2 \right\}$ [Theorem 1.1](#) yields $\Pr \left[G' \in A_{\frac{\varepsilon}{18}}^{(n-k)} \right] \geq 1 - \frac{\varepsilon}{12}$ for $G' \sim \zeta_{\mathbb{V}, \phi}^{(n-k)}$. Consequently, we have

$$\Psi_n^{(k)}(\mathbf{x}) \geq \frac{1 - \frac{\varepsilon}{12}}{1 + \frac{\varepsilon}{12}} \cdot \frac{1}{\Xi_{\mathbb{V}}(\lambda, \phi)}$$

and

$$g_{\varepsilon}(\eta) \geq \left(1 - \frac{\varepsilon}{12}\right)^2 \left(1 + \frac{\varepsilon}{12}\right)^{-1} f(\eta) \geq \left(1 - \frac{\varepsilon}{12}\right)^3 f(\eta) \geq \left(1 - \frac{\varepsilon}{4}\right) f(\eta),$$

which concludes the proof. ■

We proceed by using [Lemmas 5.3](#) and [5.5](#) to bound the total variation distance between $P_{\mathbb{V}}^{(\lambda, \phi)}$ and the output distribution of [Algorithm 2](#). However, as [Lemma 5.5](#) only provides information for point sets that

are sufficiently small compared to n , we need a different way to deal with large point configurations. To this end, the following two results are useful. The first lemma is a domination result for the size of independent sets, drawn from a hard-core model.

► **Lemma 5.6.** Let $G \in \mathcal{G}_n$ for some $n \in \mathbb{N}$ and let $\gamma \in \mathbb{R}_{\geq 0}$. For $\sigma \sim \mu_G^{(\gamma)}$ it holds that $|S_\sigma|$ is stochastically dominated by a binomial random variable with n trials and success probability $\frac{\gamma}{1+\gamma}$. ◀

Proof. We use a coupling argument to prove this statement. Consider the following procedure for sampling a set $S_n \subseteq [n]$:

1. Start with $S_0 = \emptyset$.
2. For each $i \in [n]$, set $S_i = S_{i-1} \cup \{i\}$ with probability $\Pr[\sigma(i) = 1 \mid \bigwedge_{j \in [i-1]} \sigma(j) = \mathbb{1}_{j \in S_{i-1}}]$ for $\sigma \sim \mu_G^{(\gamma)}$.

Note that the resulting set S_n follows the same distribution as S_σ for $\sigma \sim \mu_G^{(\gamma)}$. Due to the definition of this process, it suffices to consider sequences $(S_i)_{i \in [n] \cup \{0\}}$ such that the event $\{\bigwedge_{j \in [i-1]} \sigma(j) = \mathbb{1}_{j \in S_{i-1}}\}$ has non-zero probability. Further, note that

$$\Pr \left[\sigma(i) = 1 \mid \bigwedge_{j \in [i-1]} \sigma(j) = \mathbb{1}_{j \in S_{i-1}} \right] \leq \frac{\gamma}{1+\gamma}$$

for all $i \in [n]$. Now, we consider a modified process $(S'_i)_{i \in [n] \cup \{0\}}$ with $S'_0 = \emptyset$ and $S'_i = S'_{i-1} \cup \{i\}$ with probability $\frac{\gamma}{1+\gamma}$. Observe that $(S_i)_{i \in [n] \cup \{0\}}$ and $(S'_i)_{i \in [n] \cup \{0\}}$ can be coupled in such a way that $S_i \subseteq S'_i$ whenever $S_{i-1} \subseteq S'_{i-1}$ for all $i \in [n]$. As initially $S_0 = S'_0$, the same coupling yields $S_n \subseteq S'_n$. Finally, observing that $|S'_n|$ follows a binomial distribution with n trials and success probability $\frac{\gamma}{1+\gamma}$ concludes the proof. ■

The second lemma is the analog of [Lemma 5.6](#) for repulsive point processes. However, proving it is slightly more technically involved. We start by introducing some additional notation and terminology. For two counting measures $\eta_1, \eta_2 \in \mathcal{N}$, we write $\eta_1 \leq \eta_2$ if $\eta_1(B) \leq \eta_2(B)$ for every $B \in \mathcal{B}$. A measurable function $h : \mathcal{N} \rightarrow \mathbb{R}$ is called *increasing* if $h(\eta_1) \leq h(\eta_2)$ for all $\eta_1 \leq \eta_2$. Moreover, for some $\kappa \in \mathbb{R}_{\geq 0}$, let Q_κ denote the Poisson point process with intensity κ and let P be a point process that has a density f_P with respect to Q_κ . A function $\zeta : \mathcal{N} \times \mathbb{X} \rightarrow \mathbb{R}_{\geq 0}$ is called a *Papangelou intensity* for P (w.r.t. Q_κ) if, for all $\eta \in \mathcal{N}$ and $x \in \mathbb{X}$, it holds that

$$f_P(\eta + \delta_x) = \zeta(\eta, x) f_P(\eta).$$

The domination lemma we are aiming for is implied by the following result.

► **Theorem 5.7** ([\[20, Theorem 1.1\]](#)). Let Q_κ be a Poisson point process of intensity $\kappa \in \mathbb{R}_{\geq 0}$ and let P_1, P_2 be point processes that are absolutely continuous with respect to Q_κ . Assume P_1 and P_2 have Papangelou intensities ζ_1 and ζ_2 . If, for all $x \in \mathbb{X}$ and $\eta_1, \eta_2 \in \mathcal{N}$ with $\eta_1 \leq \eta_2$, $\zeta_1(\eta_1, x) \leq \zeta_2(\eta_2, x)$, then, for all increasing $h : \mathcal{N} \rightarrow \mathbb{R}$, it holds that

$$\int_{\mathcal{N}} h(\eta) P_1(d\eta) \leq \int_{\mathcal{N}} h(\eta) P_2(d\eta). \quad \blacktriangleleft$$

With that, we show the following simple domination result.

► **Lemma 5.8.** Let (\mathbb{X}, d) be a complete, separable metric space, let $\mathcal{B} = \mathcal{B}(\mathbb{X})$ be the Borel algebra and let ν be a locally finite reference measure on $(\mathbb{X}, \mathcal{B})$. Let $\mathbb{V} \subseteq \mathbb{X}$ be bounded and measurable, let $\lambda \in \mathbb{R}_{\geq 0}$ and let $\phi : \mathbb{X}^2 \rightarrow \mathbb{R}_{\geq 0} \cup \{\infty\}$ be a symmetric repulsive potential. For $\eta \sim P_{\mathbb{V}}^{(\lambda, \phi)}$ it holds that $\eta(\mathbb{V})$ is dominated by a Poisson random variable with parameter $\lambda\nu(\mathbb{V})$. ◀

Proof. Let Q_λ denote a Poisson point process with intensity λ . Note that a density of

$$f_1(\eta) = \frac{1}{\Xi_{\mathbb{V}}(\lambda, \phi)} \mathbb{1}_{\eta \in \mathcal{N}_{\mathbb{V}}} \left(\prod_{\{x, y\} \in \binom{X_\eta}{2}} e^{-N_x(\eta)N_y(\eta)\phi(x, y)} \right) \left(\prod_{x \in X_\eta} e^{-\frac{N_x(\eta)(N_x(\eta)-1)}{2}\phi(x, x)} \right) e^{\lambda\nu(\mathbb{V})}$$

is a density for $P_{\mathbb{V}}^{(\lambda, \phi)}$ with respect to Q_λ . Therefore,

$$\zeta_1(\eta, x) = \mathbb{1}_{x \in \mathbb{V}} \prod_{y \in X_\eta} e^{-N_y(\eta)\phi(x, y)}$$

is a Papangelou intensity for $P_{\mathbb{V}}^{(\lambda, \phi)}$. Moreover, let P denote the point process defined by the density $f_2(\eta) = \mathbb{1}_{\eta \in \mathcal{N}_{\mathbb{V}}}$ and observe that $\zeta_2(\eta, x) = \mathbb{1}_{x \in \mathbb{V}}$ is a Papangelou intensity for P .

For all $k \in \mathbb{N}$, let $h_k(\eta) = \mathbb{1}_{\eta(\mathbb{V}) \geq k}$ and observe that h_k is increasing. Further, note that, for all $x \in \mathbb{X}$ and $\eta_1, \eta_2 \in \mathcal{N}$, it holds that

$$\zeta_1(\eta_1, x) = \mathbb{1}_{x \in \mathbb{V}} \prod_{y \in X_{\eta_1}} e^{-N_y(\eta_1)\phi(x, y)} \leq \mathbb{1}_{x \in \mathbb{V}} = \zeta_2(\eta_2, x).$$

By [Theorem 5.7](#), this implies that for all $k \in \mathbb{N}$

$$\int_{\mathcal{N}} h_k(\eta) P_{\mathbb{V}}^{(\lambda, \phi)}(d\eta) \leq \int_{\mathcal{N}} h_k(\eta) P(d\eta).$$

Consequently, for $\eta \sim P_{\mathbb{V}}^{(\lambda, \phi)}$ and $\xi \sim P$ and for all $k \in \mathbb{N}$, it holds that

$$\Pr[\eta(\mathbb{V}) \geq k] \leq \Pr[\xi(\mathbb{V}) \geq k]$$

and observing that $\xi(\mathbb{V})$ follows a Poisson distribution with parameter $\lambda\nu(\mathbb{V})$ concludes the proof. ■

We now bound the total variation distance between the output of [Algorithm 2](#) and $P_{\mathbb{V}}^{(\lambda, \phi)}$.

► **Lemma 5.9.** Let (\mathbb{X}, d) be a complete, separable metric space, let $\mathcal{B} = \mathcal{B}(\mathbb{X})$ be the Borel algebra and let ν be a locally finite reference measure on $(\mathbb{X}, \mathcal{B})$. Let $\mathbb{V} \subseteq \mathbb{X}$ be bounded and measurable, let $\lambda \in \mathbb{R}_{\geq 0}$ and let $\phi : \mathbb{X}^2 \rightarrow \mathbb{R}_{\geq 0} \cup \{\infty\}$ be a symmetric repulsive potential. For every given $\varepsilon \in (0, 1]$, [Algorithm 2](#) is an $\frac{\varepsilon}{2}$ -approximate sampler from $P_{\mathbb{V}}^{(\lambda, \phi)}$. ◀

Proof. We start by bounding the total variation distance between $d_{\text{tv}}\left(P_{\mathbb{V}}^{(\lambda, \phi)}, \widehat{P}_\varepsilon\right)$ for \widehat{P}_ε as in [Lemma 5.3](#). The statement then follows from a coupling argument. Let Q_λ denote a Poisson point process of intensity λ . Let g_ε be the density of \widehat{P}_ε with respect to Q_λ as given in [Lemma 5.3](#) and let f denote the density of $P_{\mathbb{V}}^{(\lambda, \phi)}$ with respect to Q . Moreover, set $m = \min\left\{\sqrt{\frac{\varepsilon n}{12}}, \frac{\varepsilon}{40\lambda\nu(\mathbb{V})+\varepsilon}n\right\}$. Note that the total variation distance can be

expressed as

$$\begin{aligned} d_{\text{tv}}\left(P_{\mathbb{V}}^{(\lambda, \phi)}, \widehat{P}_\varepsilon\right) &= \int_{\mathcal{N}} |f(\eta) - g_\varepsilon(\eta)| Q_\lambda(d\eta) \\ &= \int_{\mathcal{N}} \mathbb{1}_{\eta(\mathbb{V}) \leq m} |f(\eta) - g_\varepsilon(\eta)| Q_\lambda(d\eta) + \int_{\mathcal{N}} \mathbb{1}_{\eta(\mathbb{V}) > m} |f(\eta) - g_\varepsilon(\eta)| Q_\lambda(d\eta). \end{aligned}$$

By [Lemma 5.5](#), we get

$$\int_{\mathcal{N}} \mathbb{1}_{\eta(\mathbb{V}) \leq m} |f(\eta) - g_\varepsilon(\eta)| Q_\lambda(d\eta) \leq \frac{\varepsilon}{4} \int_{\mathcal{N}} \mathbb{1}_{\eta(\mathbb{V}) \leq m} f(\eta) Q_\lambda(d\eta) \leq \frac{\varepsilon}{4}.$$

Further, it holds that

$$\begin{aligned} \int_{\mathcal{N}} \mathbb{1}_{\eta(\mathbb{V}) > m} |f(\eta) - g_\varepsilon(\eta)| Q_\lambda(d\eta) &\leq \int_{\mathcal{N}} \mathbb{1}_{\eta(\mathbb{V}) > m} f(\eta) Q_\lambda(d\eta) + \int_{\mathcal{N}} \mathbb{1}_{\eta(\mathbb{V}) > m} g_\varepsilon(\eta) Q_\lambda(d\eta) \\ &= \Pr[\xi(\mathbb{V}) > m] + \Pr\left[|Y| > m \mid G \in A_{\frac{\varepsilon}{12}}^{(n)}\right] \end{aligned}$$

for $\xi \sim P_{\mathbb{V}}^{(\lambda, \phi)}$, and G and Y as in [Algorithm 2](#).

We proceed by bounding each of these probability separately. Note that, by our choice of n it holds that $m \geq \frac{12}{\varepsilon} \lambda \nu(\mathbb{V})$. By [Lemma 5.8](#), we have $\mathbb{E}_{\xi \sim P_{\mathbb{V}}^{(\lambda, \phi)}}[\xi(\mathbb{V})] \leq \lambda \nu(\mathbb{V})$. Thus, Markov's inequality yields $\Pr[\xi(\mathbb{V}) > m] \leq \frac{\varepsilon}{12}$. Moreover, note that $|Y| = |S_\tau|$ for τ as in [Algorithm 2](#). As $\tau \sim \mu_H^{(Y(n))}$ for some $H \in A_{\frac{\varepsilon}{12}}^{(n)} \subseteq \mathcal{G}_n$ and [Lemma 5.6](#) applies to all such graphs, we get

$$\mathbb{E}\left[|Y| \mid G \in A_{\frac{\varepsilon}{12}}^{(n)}\right] \leq \frac{\gamma(n)}{1 + \gamma(n)} n \leq \gamma(n) n = \lambda \nu(\mathbb{V}).$$

Again, applying Markov's inequality gives $\Pr\left[|Y| > m \mid G \in A_{\frac{\varepsilon}{12}}^{(n)}\right] \leq \frac{\varepsilon}{12}$. Consequently, we have

$$d_{\text{tv}}\left(P_{\mathbb{V}}^{(\lambda, \phi)}, \widehat{P}_\varepsilon\right) \leq \frac{\varepsilon}{4} + \frac{\varepsilon}{6} = \frac{5}{12} \varepsilon.$$

To finish the proof, we now relate the output of [Algorithm 2](#) with \widehat{P}_ε by using a coupling argument. To this end, note that [Algorithm 2](#) can be used to sample from \widehat{P}_ε by simply restarting the sampler whenever $G \notin A_{\frac{\varepsilon}{12}}^{(n)}$. For our choice of n we know that with a probability of $\Pr\left[G \in A_{\frac{\varepsilon}{12}}^{(n)}\right] \geq 1 - \frac{\varepsilon}{12}$ only a single run of [Algorithm 2](#) is required. By this coupling, the total variation distance between the output of [Algorithm 2](#) and \widehat{P}_ε is at most $\frac{\varepsilon}{12}$. Finally, applying triangle inequality shows that the total variation distance between the output of [Algorithm 2](#) and $P_{\mathbb{V}}^{(\lambda, \phi)}$ is bounded by $\frac{5}{12} \varepsilon + \frac{\varepsilon}{12} = \frac{\varepsilon}{2}$, which concludes the proof. \blacksquare

Using [Lemma 5.9](#), we are able to prove that [Algorithm 1](#) is an ε -approximate sampler for $P_{\mathbb{V}}^{(\lambda, \phi)}$. In order to argue that [Algorithm 1](#) also satisfies the running time requirements, given in [Theorem 5.1](#), we require an efficient approximate sampler from the hard-core distribution $\mu_G^{(Y(n))}$. To this end, we use the following known result.

► **Theorem 5.10** ([\[3, Theorem 5\]](#)). Let $G = (V, E)$ be an undirected graph with maximum vertex degree

bounded by $d_G \in \mathbb{N}_{\geq 2}$ and let $\gamma \in \mathbb{R}_{\geq 0}$ with

$$\gamma < \gamma_c(d_G) = \frac{(d_G - 1)^{d_G - 1}}{(d_G - 2)^{d_G}}.$$

Then, for all $\varepsilon \in (0, 1]$, there is an ε -approximate sampler for the hard-core Gibbs distribution $\mu_G^{(\gamma)}$ with an expected running time of $O\left(|V| \ln\left(\frac{|V|}{\varepsilon}\right)\right)$. \blacktriangleleft

Proof of Theorem 5.1. We start by arguing that Algorithm 1 is an ε -approximate sampler for $P_{\mathbb{V}}^{(\lambda, \phi)}$. To this end, we show that the total variation distance between the output distributions of Algorithm 1 and Algorithm 2 is bounded by $\frac{\varepsilon}{2}$. Using the triangle inequality and Lemma 5.9 then yields the desired result. To bound the total variation distance between the Algorithm 1 and Algorithm 2 by $\frac{\varepsilon}{2}$, it suffices to construct a coupling of both algorithms such that their output coincides with a probability of at least $1 - \frac{\varepsilon}{2}$. This is, we want to find a coupling of both algorithms such that $X \neq Y$ with probability at most $\frac{\varepsilon}{2}$, where X and Y are as in Algorithm 1 and Algorithm 2.

To construct such a coupling, we start by letting both algorithms draw the same points X_1, \dots, X_n and construct the same graph G . If $d_G \geq \frac{en}{\lambda v(\mathbb{V})}$, then we may just assume $X \neq Y$. Otherwise, if $d_G < \frac{en}{\lambda v(\mathbb{V})}$, then $\sigma = \tau$ is a sufficient condition for $X = Y$. As τ is drawn from $\mu_G^{(\gamma(n))}$ and σ is drawn from an $\frac{\varepsilon}{4}$ approximation of that distribution, they can be coupled in such a way that $\Pr[\tau \neq \sigma] \leq \frac{\varepsilon}{4}$. Using this coupling of Algorithm 1 and Algorithm 2, we have

$$\Pr[X \neq Y] \leq \Pr\left[d_G \geq \frac{en}{\lambda v(\mathbb{V})}\right] + \frac{\varepsilon}{4} \cdot \Pr\left[d_G < \frac{en}{\lambda v(\mathbb{V})}\right] \leq \Pr\left[d_G \geq \frac{en}{\lambda v(\mathbb{V})}\right] + \frac{\varepsilon}{4}.$$

Therefore, it remains to prove that $d_G \geq \frac{en}{\lambda v(\mathbb{V})}$ with probability at most $\frac{\varepsilon}{4}$, where $G \sim \zeta_{\mathbb{V}, \phi}^{(n)}$. We follow a similar arguments as in the proof of Theorem 1.2. Note that, for our choice of n , there exists $z \geq 3 \ln\left(\frac{4e}{\varepsilon}\right) \max\left\{\frac{1}{e - \lambda C_\phi}, \frac{\lambda C_\phi}{(e - \lambda C_\phi)^2}\right\} \lambda v(\mathbb{V})$ such that $n = 2z \ln(z)^2$. Moreover, we have $n \geq e \geq 2$ and, for $v(\mathbb{V})$ (consequently z) sufficiently large, it holds that $2z \ln(z)^2 \geq 2z \ln(2z \ln(z)^2) = 2z \ln(n)$. Therefore, we have

$$\begin{aligned} n - 1 &\geq \frac{n}{2} \\ &\geq 3 \ln\left(\frac{4e}{\varepsilon}\right) \max\left\{\frac{1}{e - \lambda C_\phi}, \frac{\lambda C_\phi}{(e - \lambda C_\phi)^2}\right\} \lambda v(\mathbb{V}) \ln(n) \\ &\geq 3 \ln\left(\frac{4n}{\varepsilon}\right) \max\left\{\frac{1}{e - \lambda C_\phi}, \frac{\lambda C_\phi}{(e - \lambda C_\phi)^2}\right\} \lambda v(\mathbb{V}) \end{aligned}$$

and by Lemma 4.4

$$\Pr\left[d_G \geq \frac{en}{\lambda v(\mathbb{V})}\right] \leq \Pr\left[d_G \geq \left(1 + \frac{e - \lambda C_\phi}{\lambda C_\phi}\right) \frac{n - 1}{v(\mathbb{V})} C_\phi\right] \leq \frac{\varepsilon}{4}.$$

To prove Theorem 5.1, it remains to show that Algorithm 1 satisfies the given running time requirements. To this end, note that, for all $\lambda < \frac{e}{C_\phi}$, it holds that $n \in \tilde{O}(v(\mathbb{V})^2 \varepsilon^{-3})$. Therefore, sampling X_1, \dots, X_n requires a running time of $\tilde{O}(nt_{\mathbb{V}}) = \tilde{O}(v(\mathbb{V})^2 \varepsilon^{-3} t_{\mathbb{V}})$. Moreover, the graph can be constructed in time $\tilde{O}(n^2 t_\phi) = \tilde{O}(v(\mathbb{V})^4 \varepsilon^{-6} t_\phi)$ and $d_G \geq \frac{en}{\lambda v(\mathbb{V})}$ can be checked in $O(1)$ if we keep track of d_G while constructing

the graph. Finally, for $d_G < \frac{en}{\lambda v(\mathbb{V})}$ it holds that

$$\gamma(n) \leq \frac{\lambda v(\mathbb{V})}{n} < \frac{e}{d_G} < \gamma_c(d_G).$$

Thus, [Theorem 5.10](#) guarantees the existence of an $\frac{\varepsilon}{8}$ -approximate sampler from $\mu_G^{(\gamma(n))}$ with an expected running time in $O(n \ln(\frac{n}{\varepsilon})) = O\left(v(\mathbb{V})^2 \varepsilon^{-3} \ln\left(\frac{v(\mathbb{V})}{\varepsilon}\right)\right)$. Note that, by Markov's inequality, the probability that this sampler takes more than $\frac{8}{\varepsilon}$ times its expected running time is bounded by $\frac{\varepsilon}{8}$. Therefore, if we run the sampler from [Theorem 5.10](#) with an error bound of $\frac{\varepsilon}{8}$ and, whenever the algorithm takes more than $\frac{8}{\varepsilon}$ times its expected running time, stop it and return an arbitrary spin configuration, this results in an $\frac{\varepsilon}{4}$ -approximate sampler with a guaranteed running time in $\tilde{O}(v(\mathbb{V})^2 \varepsilon^{-4})$. Consequently, [Algorithm 1](#) runs in time $\tilde{O}(v(\mathbb{V})^2 \varepsilon^{-4} + v(\mathbb{V})^2 \varepsilon^{-3} t_{\mathbb{V}} + v(\mathbb{V})^4 \varepsilon^{-6} t_{\phi})$, which concludes the proof. \blacksquare

6 Potential-weighted connective constant and strong spatial mixing

Throughout this section, we consider the setting introduced in [Section 4](#). In this section, we relate the potential-weighted connective constant Δ_{ϕ} of a repulsive potential with a high-probability bound on a modified version of the connective constant of a graph $G \sim \zeta_{\mathbb{V}, \phi}^{(n)}$. This modified connective constant represents the growth rate of a truncated version of the self-avoiding walk tree as used by Weitz [\[47\]](#). Besides giving a graphical interpretation for the potential-weighted connective constant, an immediate consequence of the result is that the hard-core models studied in [Section 4](#) with high probability exhibit strong spatial mixing for $\gamma < \frac{e}{\Delta_{\phi}}$.

6.1 Potential-weighted connective constant

The potential-weighted connective constant was introduced in [\[33\]](#) to measure the strength of interaction induced by a potential ϕ in a way that is, compared to the temperedness constant C_{ϕ} , more sensitive to the particular geometry of the underlying space \mathbb{X} . To this end, we set $V_{\phi}(0) = 1$ and, for $k \in \mathbb{N}_{\geq 1}$,

$$V_{\phi}(k) = \sup_{x_0 \in \mathbb{X}} \int_{\mathbb{X}^k} \prod_{j=1}^k \left(\exp\left(-\sum_{i=0}^{j-2} \mathbb{1}_{d(x_i, x_j) < d(x_i, x_{i+1})} \phi(x_i, x_j)\right) \cdot \left(1 - e^{-\phi(x_{j-1}, x_j)}\right) \right) \nu^k(d\mathbf{x}),$$

where $\mathbf{x} = (x_1, \dots, x_k)$. The potential-weighted connective constant is now defined as

$$\Delta_{\phi} = \lim_{k \rightarrow \infty} V_{\phi}(k)^{1/k} = \inf_{k \rightarrow \infty} V_{\phi}(k)^{1/k},$$

where existence of the limit and the second equality are implied by the fact that we assume ϕ to be repulsive, which implies that $V_{\phi}(k)$ is sub-multiplicative. Note that for all $k \in \mathbb{N}$ it holds that $V_{\phi}(k) \leq C_{\phi}^k$ and therefore $\Delta_{\phi} \leq C_{\phi}$.

6.2 Strong spatial mixing

Strong spatial mixing is a frequently used notion of correlation decay in discrete spin systems. In this section, we focus on strong spatial mixing for the hard-core model. Recall that for a graph $G = (V_G, E_G)$ and a parameter $\gamma \in \mathbb{R}_{\geq 0}$ we write $\mu_G^{(\gamma)}$ for the hard-core distribution on G at weight γ , which is a distribution

on Σ_G the set of all functions $\Sigma_G = \{\sigma : V \rightarrow \{0, 1\}\}$. We extend this notation to conditional distributions. To this end, let $S \subseteq V_G$ and let $\tau : S \rightarrow \{0, 1\}$. Write $\sigma_S = \tau$ for the event that $\sigma \sim \mu_G^{(\gamma)}$ coincides with τ on S . We call τ *feasible* if $\mu_G^{(\gamma)}(\sigma_S = \tau) > 0$. In that case, we write

$$\mu_G^{(\gamma)}(\cdot \mid \sigma_S = \tau) = \frac{\mu_G^{(\gamma)}(\cdot)}{\mu_G^{(\gamma)}(\sigma_S = \tau)}$$

for the distribution of $\sigma \sim \mu_G^{(\gamma)}$ conditioned on $\sigma_S = \tau$. Often, strong spatial mixing is phrased in terms of the so-called *occupation ratios*. For G and γ as above and $v \in V_G$ we write

$$R_G^{(\gamma)}(v) = \frac{\mu_G^{(\gamma)}(\sigma(v) = 1)}{\mu_G^{(\gamma)}(\sigma(v) = 0)}.$$

Further, for $S \subset V_G$ with $v \notin S$ and feasible $\tau : S \rightarrow \{0, 1\}$ we define $R_G^{(\gamma)}(v \mid \tau)$ analogously using the distribution $\mu_G^{(\gamma)}(\cdot \mid \sigma_S = \tau)$ instead. Based on that, we now state the definition of strong spatial mixing as given in [40].

► **Definition 6.1** ([40]). The hard-core model with vertex activity $\gamma \in \mathbb{R}_{\geq 0}$ is set to satisfy strong spatial mixing on a family of graphs \mathcal{F} if there exists a constant $\delta \in [0, 1)$ such that for all $G \in \mathcal{F}$, vertices $v \in V_G$, $S \subseteq V_G \setminus \{v\}$, and feasible $\tau, \tau' : S \rightarrow \{0, 1\}$ it holds that

$$\left| R_G^{(\gamma)}(v \mid \tau) - R_G^{(\gamma)}(v \mid \tau') \right| \leq O(\delta^s),$$

where s is the graph distance between v and the vertices on which τ and τ' differ (i.e., $\{u \in S \mid \tau(u) \neq \tau'(u)\}$). We call δ the decay rate. ◀

► **Remark 6.2.** Note that this definition is actually weaker than the definition of (exponential) strong spatial mixing that is usually used in the literature (cf. [47] Definition 2.2 and the remark following it). Usually, this definition requires the existence of constants $\alpha \geq 0, \beta > 0$ (independent of G, v, τ_1 and τ_2) such that

$$\left| R_G^{(\gamma)}(v \mid \tau) - R_G^{(\gamma)}(v \mid \tau') \right| \leq \alpha e^{-\beta s}.$$

In contrast, Definition 6.1 only requires such a decay if the distance $s \geq s_0$ is sufficiently large, where the required lower bound s_0 might depend on $|V_G|$. In fact, this is the case for the strong spatial mixing result for families of finite graphs in [40], which requires $s \geq s_0 \in \Theta(\ln(|V_G|))$. Such a bound on s_0 still allows for applying Weitz's algorithm to obtain an efficient deterministic approximation algorithm for the partition function of the hard-core model. However, it is not sufficient for other applications that require the exponential decay to also hold at constant distances, such as recent exact sampling algorithms [2, 16] or the rapid mixing result for Glauber dynamics in [47]. ◀

6.3 Self-avoiding walk tree and Weitz tree

Due to [47], it is well known that strong spatial mixing on a given graph can be studied in terms of a tree construction that is closely related to the *self-avoiding walk tree*. We refer to this construction as the *Weitz tree*. Moreover, it was shown by [40] that bounding the growth-rate of Weitz trees for a family of graphs can be used to find a vertex activity regime in which the graph family exhibits strong spatial mixing. In the following paragraphs, we briefly introduce both trees.

Self-avoiding walk tree Given an undirected graph $G = (V_G, E_G)$ and a vertex $r \in V_G$, we denote by $T_{r,G} = (V_T, E_T)$ the self-avoiding walk tree with root r . It is constructed as follow:

1. Let $\mathcal{P}_{r,G}$ denote the set of all simple paths $p = r, v_1, \dots, v_k$ (identified by their vertex sequence) in G starting at r , where we call k the length of p . V_T contains exactly one vertex w_p for every such path $p \in \mathcal{P}_{r,G}$. In particular, there is a unique vertex that corresponds to the path $p = r$ of length 0, which is denoted by w_r and considered the root of $T_{r,G}$.
2. For two paths $p, p' \in \mathcal{P}_{r,G}$ we say that $w_p \in V_T$ is a child of $w_{p'} \in V_T$ (and connect them by an edge in E_T) if and only if p is obtained from p' by adding one vertex. That is, if $p' = r, v_1, \dots, v_k$, then $p = r, v_1, \dots, v_k, v_{k+1}$ for some vertex $v_{k+1} \notin V_G \setminus \{r, v_1, \dots, v_k\}$ adjacent to v_k .

Weitz tree For our purpose, it is crucial to differentiate between the self-avoiding walk tree and a truncated version that we call the Weitz tree. Its construction is analogous to the self-avoiding walk tree above but with an additional restriction on the set of paths involved. This restriction on the paths allows for some degrees of freedom, as it depends on assigning an ordering to the neighbors $N_G(v)$ of each vertex $v \in V_G$. Formally, we call a family of functions $F = (f_v)_{v \in V_G}$ a *neighborhood ordering* for G if, for every $v \in V_G$, it holds that f_v is a bijection $N_G(v) \rightarrow [|N_G(v)|]$. For a vertex $r \in V_G$ we now write $\mathcal{P}_{r,G}^F$ for the set of simple paths $p = v_0, v_1, \dots, v_k$ in G with the following properties

1. It holds that $v_0 = r$.
2. For every $2 \leq i \leq k$ and all $0 \leq j \leq i - 2$ it holds that, if $v_i \in N_G(v_j)$, then $f_{v_j}(v_i) > f_{v_j}(v_{j+1})$.

Obviously it holds that $\mathcal{P}_{r,G}^F \subseteq \mathcal{P}_{r,G}$. The Weitz tree $T_{r,G}^F$ with root r is now defined analogously to the self-avoiding walk tree $T_{r,G}$ but restricted to paths in $\mathcal{P}_{r,G}^F$. It is not hard to see that $T_{r,G}^F$ is in fact a subtree of $T_{r,G}$. Since we are going to study a notion of growth rate of $T_{r,G}^F$, it is useful to denote by $L_{r,G}^F(k)$ the number of vertices at layer $k \in \mathbb{N}$ of $T_{r,G}^F$. This is equal to the number of paths $p = r, v_1, \dots, v_k \in \mathcal{P}_{r,G}^F$ of length k .

► **Remark 6.3.** Note that this construction of the Weitz tree is not exactly as described in [47]. The goal of constructing those trees is to study hard-core models on them. In the original construction, paths were allowed to close cycles. Whenever this happens, the spin of the vertex that closes the cycle was fixed to either 0 or 1, depending on the chosen neighborhood ordering. For the hard-core model, this corresponds to either removing the vertex (for spin 0), or the vertex and all its neighbors (for spin 1). This procedure leads to our notion of the Weitz tree. ◀

6.4 Connective constant and strong spatial mixing

With the definition of strong spatial mixing and the construction of the Weitz tree given, we now get to the definition of the connective constant and its implications for strong spatial mixing. Our definition of connective constant is inspired by [40]. However, there are two things that should be noted. Firstly, we emphasize that our definition of connective constant refers to the Weitz tree instead of the full self-avoiding walk. Secondly, for our application to random graphs from $\zeta_{\mathbb{V},\phi}^{(n)}$, it is important which neighborhood ordering is used for constructing the tree. To reflect this, we use the following definition.

► **Definition 6.4.** We say a family of graphs \mathcal{F} has a connective constant bounded by Δ if there are constants $a, c \geq 0$ such that for all $G \in \mathcal{F}$ the following holds: there is a neighborhood ordering F for G such that for all $m \geq a \ln(|V_G|)$ and all $r \in V_G$ we have

$$\sum_{k=0}^m L_{r,G}^F(k) \leq c\Delta^m. \quad \blacktriangleleft$$

In [40], it was proven that a bound on the connective constant of a graph family immediately translates to a regime of strong spatial mixing. For our purposes, we will actually need some more detailed information about the rate of decay. The following statement can be extracted from the proof of the main theorem in [40].

► **Theorem 6.5 ([40]).** Suppose \mathcal{F} is a family of graphs with connective constant bounded by Δ as in Definition 6.4 for constants a and c . For all $\varepsilon > 0$, $\gamma \leq e^{-\varepsilon}\gamma_c(\Delta)$, and graphs $G \in \mathcal{F}$ the following holds: For all $v \in V_G$, all $S \subseteq V_G \setminus \{v\}$, and all feasible $\tau, \tau' : S \rightarrow \{0, 1\}$ that only differ at distance $s \geq a \ln(|V_G|)$ from v , it holds that

$$\left| R_G^{(\gamma)}(v \mid \tau) - R_G^{(\gamma)}(v \mid \tau') \right| \leq c^{1/q} \frac{M}{L} e^{-\varepsilon s/q},$$

where $1 < q \leq 2$, $M = \sinh^{-1}(\sqrt{\gamma})$ and $L = \frac{1}{2\sqrt{\gamma \cdot (\gamma+1)}}$. ◀

► **Remark 6.6.** Theorem 6.5 follows immediately from tracking the constants in the proof of the main theorem of [41]. Note further that in the proof in [41] the bound from Theorem 6.5 is stated with $\frac{L}{M}$ instead of $\frac{M}{L}$. However, closely inspecting the proof and the lemmas used therein reveals that this is attributed to a typo. ◀

6.5 Connective constant and spatial mixing for random discretizations of Gibbs point processes

Our first main result of this section is the following bound on the connective constant for random graphs from $\zeta_{\mathbb{V},\phi}^{(n)}$.

► **Theorem 6.7.** Let $\mathbb{V} \subseteq \mathbb{X}$ be a bounded and measurable region with volume $\nu(\mathbb{V}) > 0$. Let $\phi : \mathbb{X}^2 \rightarrow \mathbb{R}_{\geq 0} \cup \{\infty\}$ be a symmetric repulsive potential with $C_\phi < \infty$ and $\Delta_\phi > 0$. For every $\varepsilon > 0$ there exists some $n_0 \in \Theta(\nu(\mathbb{V}))$ such that for all $n \geq n_0$ the following holds: There exists a family of graphs $\mathcal{F} \subseteq \mathcal{G}_n$ with connective constant bounded by $e^\varepsilon \frac{n}{\nu(\mathbb{V})} \Delta_\phi$ such that the constants a, c from Definition 6.4 are independent of \mathbb{V} and n , and for $G \sim \zeta_{\mathbb{V},\phi}^{(n)}$ it holds that

$$\Pr[G \in \mathcal{F}] \geq 1 - \frac{1}{n}. \quad \blacktriangleleft$$

Proof. Fix some $\varepsilon > 0$. By the definition of Δ_ϕ and the fact that $\Delta_\phi > 0$, we know that there is some k_0 such that $V_\phi(k) \leq e^{k\varepsilon/2} \Delta_\phi^k$ for all $k \geq k_0$. We set $a = 4\varepsilon^{-1}$, $c = (1 - e^{-\varepsilon/2}) \cdot ((C_\phi/\Delta_\phi)^{k_0} + 2)$, and $n_0 = \max\left\{\frac{2}{\Delta_\phi} \nu(\mathbb{V}), e^{a^{-1}k_0}\right\}$.

Next, fix any $n \geq n_0$. Our goal is to show that, for $G \sim \zeta_{\mathbb{V},\phi}^{(n)}$, we can find some neighborhood ordering for G such that, with probability at least $1 - 1/n$, the requirements of Definition 6.4 are satisfied for a, c as above and $\Delta = e^\varepsilon \frac{n}{\nu(\mathbb{V})} \Delta_\phi$.

We start by constructing the class of neighborhood orderings we consider. Let $G \in \mathcal{G}_n$ be a graph on $[n]$. For every $\mathbf{x} = (x_1, \dots, x_n) \in \mathbb{V}^n$ we construct a neighborhood ordering $F_{\mathbf{x}, G} = (f_i)_{i \in [n]}$ on G such that, for every vertex $i \in [n]$ and neighbors $j_1, j_2 \in N_G(i)$, it holds that $d(x_i, x_{j_1}) < d(x_i, x_{j_2})$ implies $f_i(j_1) < f_i(j_2)$. That is, we order the neighbors of each vertex increasingly by distance, breaking ties arbitrarily (e.g., by vertex IDs). It now suffices to show that, for $G \sim \zeta_{\mathbb{V}, \phi}^{(n)}$, the following event has probability at least $1 - 1/n$: there is a sequence of points $\mathbf{x} \in \mathbb{V}^n$ such that for all $m \geq a \ln(n)$ and all $r \in [n]$

$$\sum_{k=0}^m L_{r, G}^{F_{\mathbf{x}, G}}(k) \leq c \left(e^\varepsilon \frac{n}{v(\mathbb{V})} \Delta_\phi \right)^m.$$

We denote this event by A , so that the desired statement is simply expressed as $\zeta_{\mathbb{V}, \phi}^{(n)}(A) \geq 1 - 1/n$.

To prove this, we study a distribution κ on the space $\mathbb{V}^n \times \mathcal{G}_n$, equipped with the product sigma field $\mathcal{B}_{\mathbb{V}}^n \otimes 2^{\mathcal{G}_n}$. Consider the following procedure:

1. For each $i \in [n]$, draw a uniform random point $X_i \sim u_{\mathbb{V}}$ independently. We call the resulting random vector $X = (X_i)_{i \in [n]}$
2. Construct a graph G on vertex set $[n]$ as follows. For all $i, j \in [n]$ with $i \neq j$, connect vertices i and j with an edge with probability $1 - e^{-\phi(X_i, X_j)}$ independently.

We take κ to be the distribution of (X, G) generated as above. We further write κ_X and κ_G for the respective marginals and note that $\kappa_X = u_{\mathbb{V}}^n$ and $\kappa_G = \zeta_{\mathbb{V}, \phi}^{(n)}$. Further, let B denote the following event: for all $m \geq a \ln(n)$ and all $r \in [n]$ it holds that

$$\sum_{k=0}^m L_{r, G}^{F_{X, G}}(k) \leq c \left(e^\varepsilon \frac{n}{v(\mathbb{V})} \Delta_\phi \right)^m.$$

Note that the main difference between the events A and B is that A asks for the existence of a point sequence \mathbf{x} for a given random graph G , whereas B is a statement about a random pair (X, G) . In particular, it holds that $B \subseteq \mathbb{V}^n \times A$ and consequently

$$\zeta_{\mathbb{V}, \phi}^{(n)}(A) = \kappa_G(A) = \kappa(\mathbb{V}^n \times A) \geq \kappa(B).$$

Thus, the theorem is proven by showing that $\kappa(B) \geq 1 - 1/n$.

We prove this by first bounding the expectation of $L_{r, G}^{F_{X, G}}(k)$ with respect to κ for every fixed $r \in [n]$. To this end, define $\psi : \mathbb{X}^{k+1} \rightarrow [0, 1]$ by

$$\psi(y_0, y_1, \dots, y_k) = \prod_{j=1}^k \left(\exp \left(- \sum_{i=0}^{j-2} \mathbb{1}_{d(y_i, y_j) < d(y_i, y_{i+1})} \phi(y_i, y_j) \right) \cdot \left(1 - e^{-\phi(y_{j-1}, y_j)} \right) \right).$$

For every sequence of distinct vertices $i_1, \dots, i_k \in [n] \setminus \{r\}$ it now holds that

$$\mathbb{E} \left[\mathbb{1}_{r, i_1, \dots, i_k \in \mathcal{P}_{r, G}^{F_{X, G}}} \mid X \right] \leq \psi(X_r, X_{i_1}, \dots, X_{i_k}).$$

Applying linearity of expectation to sum over all such sequences of distinct vertices $i_1, \dots, i_k \in [n] \setminus \{r\}$

and applying law of total expectation then yields that

$$\begin{aligned} \mathbb{E} \left[L_{r,G}^{F_{X,G}}(k) \right] &\leq \nu(\mathbb{V})^{-n} \int_{\mathbb{V}^n} \sum_{i_1, \dots, i_k} \psi(x_r, x_{i_1}, \dots, x_{i_k}) \nu^n(d\mathbf{x}) \\ &\leq \nu(\mathbb{V})^{-k} \frac{(n-1)!}{(n-1-k)!} \sup_{z_0 \in \mathbb{V}} \int_{\mathbb{V}^k} \psi(z_0, z_1, \dots, z_k) \nu^k(d\mathbf{z}) \\ &\leq \left(\frac{n}{\nu(\mathbb{V})} \right)^k V_\phi(k), \end{aligned}$$

where $\mathbf{x} = (x_1, \dots, x_n)$ and $\mathbf{z} = (z_1, \dots, z_k)$.

Next, we aim to obtain a tail bound for $\sum_{k=0}^m L_{r,G}^{F_{X,G}}(k)$ for every $m \geq a \ln(n)$. To this end, we first bound the expectation, starting with splitting up the sum as

$$\mathbb{E} \left[\sum_{k=0}^m L_{r,G}^{F_{X,G}}(k) \right] \leq \sum_{k=0}^m \left(\frac{n}{\nu(\mathbb{V})} \right)^k V_\phi(k) = \sum_{k=0}^{k_0-1} \left(\frac{n}{\nu(\mathbb{V})} \right)^k V_\phi(k) + \sum_{k=k_0}^m \left(\frac{n}{\nu(\mathbb{V})} \right)^k V_\phi(k).$$

We proceed by bounding each of the sums separately. For the first sum, note that it trivially holds that $V_\phi(k) \leq C_\phi^k$ and, in particular, $\Delta_\phi \leq C_\phi$. Moreover, for our choice of n_0 , we $\frac{n}{\nu(\mathbb{V})} C_\phi \geq 2$. Combining both observations yields

$$\sum_{k=0}^{k_0-1} \left(\frac{n}{\nu(\mathbb{V})} \right)^k V_\phi(k) \leq \sum_{k=0}^{k_0-1} \left(\frac{n}{\nu(\mathbb{V})} C_\phi \right)^k = \frac{\left(\frac{n}{\nu(\mathbb{V})} C_\phi \right)^{k_0} - 1}{\frac{n}{\nu(\mathbb{V})} C_\phi - 1} \leq \left(\frac{n}{\nu(\mathbb{V})} C_\phi \right)^{k_0}.$$

For the second sum, recall that $V_\phi(k) \leq e^{k\epsilon/2} \Delta_\phi^k$ for all $k \geq k_0$. Therefore, we have

$$\sum_{k=k_0}^m \left(\frac{n}{\nu(\mathbb{V})} \right)^k V_\phi(k) \leq \sum_{k=k_0}^m \left(e^{\epsilon/2} \frac{n}{\nu(\mathbb{V})} \Delta_\phi \right)^k \leq \frac{\left(e^{\epsilon/2} \frac{n}{\nu(\mathbb{V})} \Delta_\phi \right)^{m+1}}{e^{\epsilon/2} \frac{n}{\nu(\mathbb{V})} \Delta_\phi - 1}.$$

Combining both bounds yields

$$\mathbb{E} \left[\sum_{k=0}^m L_{r,G}^{F_{X,G}}(k) \right] \leq \left(\frac{\left(\frac{n}{\nu(\mathbb{V})} C_\phi \right)^{k_0}}{\left(e^{\epsilon/2} \frac{n}{\nu(\mathbb{V})} \Delta_\phi \right)^m} + \frac{e^{\epsilon/2} \frac{n}{\nu(\mathbb{V})} \Delta_\phi}{e^{\epsilon/2} \frac{n}{\nu(\mathbb{V})} \Delta_\phi - 1} \right) \cdot \left(e^{\epsilon/2} \frac{n}{\nu(\mathbb{V})} \Delta_\phi \right)^m.$$

Since $e^{\epsilon/2} \frac{n}{\nu(\mathbb{V})} \Delta_\phi \geq 2$ for $n \geq n_0$, we further bound

$$\begin{aligned} \mathbb{E} \left[\sum_{k=0}^m L_{r,G}^{F_{X,G}}(k) \right] &\leq \left(\left(\frac{\frac{n}{\nu(\mathbb{V})} C_\phi}{e^{\epsilon/2} \frac{n}{\nu(\mathbb{V})} \Delta_\phi} \right)^{k_0} + 2 \right) \cdot \left(e^{\epsilon/2} \frac{n}{\nu(\mathbb{V})} \Delta_\phi \right)^m \\ &\leq \left(\left(\frac{C_\phi}{\Delta_\phi} \right)^{k_0} + 2 \right) \cdot \left(e^{\epsilon/2} \frac{n}{\nu(\mathbb{V})} \Delta_\phi \right)^m. \end{aligned}$$

Using this upper bound for the expectation and Markov's inequality, we have for $(X, G) \sim \kappa$ that

$$\Pr \left[\sum_{k=0}^m L_{r,G}^{F_{X,G}}(k) \geq c \left(e^\varepsilon \frac{n}{\nu(\mathbb{V})} \Delta_\phi \right)^m \right] \leq (1 - e^{-\varepsilon/2}) \cdot e^{-m\varepsilon/2}.$$

Applying the union bound over $m \geq m_0 = a \ln(n)$ we have

$$\Pr \left[\exists m \geq m_0 : \sum_{k=0}^m L_{r,G}^{F_{X,G}}(k) \geq c \left(e^\varepsilon \frac{n}{\nu(\mathbb{V})} \Delta_\phi \right)^m \right] \leq (1 - e^{-\varepsilon/2}) \cdot \sum_{m=m_0}^{\infty} e^{-m\varepsilon/2} \leq n^{-2},$$

where the last inequality is obtained by factoring out $e^{-m_0\varepsilon/2}$ and using the fact that $m_0 = a \ln(n) \geq 4\varepsilon^{-1} \ln(n)$. Finally, applying the union bound over the choice of root vertices r shows that $\kappa(B^c) \leq 1/n$ and consequently $\kappa(B) \geq 1 - 1/n$, proving the theorem. \blacksquare

We proceed by studying which notion of strong spatial mixing for graphs from $\zeta_{\mathbb{V},\phi}^{(n)}$ we can obtain from [Theorem 6.7](#). To this end, we apply [Theorem 6.5](#), which yields the following result.

► Theorem 6.8. Suppose $\mathbb{V} \subseteq \mathbb{X}$ is a bounded and measurable region with volume $\nu(\mathbb{V}) > 0$. Let $\phi : \mathbb{X}^2 \rightarrow \mathbb{R}_{\geq 0} \cup \{\infty\}$ be a symmetric repulsive potential with $C_\phi < \infty$ and $\Delta_\phi > 0$, and let $\lambda < \frac{e}{\Delta_\phi}$. There exists some $n_0 \in \Theta(\nu(\mathbb{V}))$ such that for all $n \geq n_0$ there is a family of graphs $\mathcal{F} \subseteq \mathcal{G}_n$ with the following properties:

1. For $G \sim \zeta_{\mathbb{V},\phi}^{(n)}$ it holds that

$$\Pr[G \in \mathcal{F}] \geq 1 - \frac{1}{n}.$$

2. Set $\gamma = \lambda \cdot \frac{\nu(\mathbb{V})}{n}$. There are constants $\alpha \geq 0, \beta > 0, a \geq 0$ independent n and \mathbb{V} such that for every $G \in \mathcal{F}$, all $v \in V_G, S \subseteq V_G \setminus \{v\}$ and feasible $\tau, \tau' : S \rightarrow \{0, 1\}$ that only differ at distance $s \geq a \ln(n)$ from v it holds that

$$\left| R_G^{(\gamma)}(v \mid \tau) - R_G^{(\gamma)}(v \mid \tau') \right| \leq \alpha e^{-\beta s}. \quad \blacktriangleleft$$

Proof. First, note that for $\lambda < \frac{e}{\Delta_\phi}$ we can choose $\varepsilon > 0$ such that $\lambda \leq \frac{e^{1-\varepsilon}}{\Delta_\phi}$. By [Theorem 6.7](#) we know that for $n_0 \in \Theta(\nu(\mathbb{V}))$ sufficiently large it holds that for all $n \geq n_0$ that there is a graph family $\mathcal{F} \subseteq \mathcal{G}_n$ with connective constant bounded by $\Delta = e^{\varepsilon/2} \frac{n}{\nu(\mathbb{V})} \Delta_\phi$ such that $G \sim \zeta_{\mathbb{V},\phi}^{(n)}$ is in \mathcal{F} with probability at least $1 - \frac{1}{n}$. Thus, \mathcal{F} satisfies the first requirement of our theorem.

We proceed by establishing the second part of the theorem. To this end, let \mathcal{F} and Δ be as above. Further, let a, c be the constants given by [Theorem 6.7](#) and recall that they are independent of n and $\nu(\mathbb{V})$. Observe that

$$\gamma = \lambda \cdot \frac{\nu(\mathbb{V})}{n} \leq \frac{e^{1-\varepsilon}}{\Delta_\phi} \cdot \frac{\nu(\mathbb{V})}{n} \leq \frac{e^{1-\varepsilon/2}}{\Delta}.$$

Thus, applying [Theorem 6.5](#) proves our claim as soon as we show that $\frac{M}{L}$ for $M = \sinh^{-1}(\sqrt{\gamma})$ and $L = \frac{1}{2\sqrt{\gamma \cdot (\gamma+1)}}$ is uniformly bounded in n and $\nu(\mathbb{V})$. Recalling that $\gamma = \lambda \cdot \frac{\nu(\mathbb{V})}{n}$, we see that $\frac{M}{L}$ is decreasing in n and bounded by a constant independent of $\nu(\mathbb{V})$ as soon as $n_0 \geq \lambda \cdot \nu(\mathbb{V})$. \blacksquare

► Remark 6.9. Note that from the proof of [Theorem 6.8](#) it actually even follows that the parameter α goes to 0 as n increases. This is to be expected, since the occupation ratios go to zero as $\gamma(n)$ decreases. \blacktriangleleft

We obtain the following corollary.

► **Corollary 6.10.** Suppose $\mathbb{V} \subseteq \mathbb{X}$ is a bounded and measurable region with volume $\nu(\mathbb{V}) > 0$. Let $\phi : \mathbb{X}^2 \rightarrow \mathbb{R}_{\geq 0} \cup \{\infty\}$ be a symmetric repulsive potential with $C_\phi < \infty$ and $\Delta_\phi > 0$, and let $\lambda < \frac{e}{\Delta_\phi}$. There exists some $n_0 \in \Theta(\nu(\mathbb{V}))$ such that for all $n \geq n_0$ there is a family of graphs $\mathcal{F} \subseteq \mathcal{G}_n$ with the following properties:

1. For $G \sim \zeta_{\mathbb{V}, \phi}^{(n)}$ it holds that

$$\Pr[G \in \mathcal{F}] \geq 1 - \frac{1}{n}.$$

2. The hard-core model with activity $\gamma = \lambda \cdot \frac{\nu(\mathbb{V})}{n}$ exhibits strong spatial mixing in the sense of [Definition 6.1](#) on \mathcal{F} with decay rate independent of n and \mathbb{V} . ◀

We finish this section with discussing some algorithmic consequences of the results above. However, we keep the discussion informal since the results can be obtained from standard techniques. We start with the observation that [Theorem 6.5](#) does in fact not only imply strong spatial mixing on each graph in the family \mathcal{F} , but it also implies a notion of strong spatial mixing in the Weitz trees that were used for bounding the connective constant of \mathcal{F} . More precisely, each of the trees exhibits strong spatial mixing in term of the occupation ratio of the root, given boundary conditions sufficiently far down the tree (see [\[40\]](#)). This can be turned into a deterministic approximation algorithm for the partition function of the hard-core model as argued in [\[47, 40\]](#) with running time $|V_G|^{\mathcal{O}(\log(\Delta))}$ for every graph in $G \in \mathcal{F}$. Using [Theorem 1.1](#), we may use this to obtain a randomized algorithm with running time $\nu(\mathbb{V})^{\mathcal{O}(\log(\nu(\mathbb{V})))}$ (i.e., quasi-polynomial in $\nu(\mathbb{V})$) for approximating $\Xi_{\mathbb{V}}(\lambda, \phi)$ for $\lambda < \frac{e}{\Delta_\phi}$ via the following procedure: We first draw a random graph $G \sim \zeta_{\mathbb{V}, \phi}^{(n)}$ together with its (random) vertex locations $X = X_1, \dots, X_n \in \mathbb{V}^n$ for n sufficiently large to satisfy [Theorem 1.1](#) and [Theorem 6.8](#). We then use the algorithm given in [\[47\]](#) together with the Weitz trees based on the neighborhood ordering $F_{X, G}$ (see proof of [Theorem 6.7](#)) to approximate the hard-core partition function on G at vertex activity $\lambda \frac{\nu(\mathbb{V})}{n}$. An analogous procedure can be used to approximately sample from $P_{\mathbb{V}}^{(\lambda, \phi)}$ with similar running time.

Acknowledgements

Andreas Göbel was funded by the project PAGES (project No. 467516565) of the German Research Foundation (DFG).

References

- [1] Konrad Anand, Andreas Göbel, Marcus Pappik, and Will Perkins. Perfect sampling for hard spheres from strong spatial mixing. In *Random*, 2023.
- [2] Konrad Anand and Mark Jerrum. Perfect sampling in infinite spin systems via strong spatial mixing. *SIAM Journal on Computing*, 51(4):1280–1295, 2022.
- [3] Nima Anari, Vishesh Jain, Frederic Koehler, Huy Tuan Pham, and Thuy-Duong Vuong. Entropic independence II: optimal sampling and concentration via restricted modified log-sobolev inequalities. *CoRR*, abs/2111.03247, 2021. URL: <https://arxiv.org/abs/2111.03247>.

- [4] Nima Anari, Kuikui Liu, and Shayan Oveis Gharan. Spectral independence in high-dimensional expanders and applications to the hardcore model. *SIAM Journal on Computing*, 0(0):FOCS20–1–FOCS20–37, 2021. doi:10.1137/20M1367696.
- [5] Andreas K. Bacher, Thomas B. Schröder, and Jeppe C. Dyre. Explaining why simple liquids are quasi-universal. *Nature Communications*, 5(1):1–7, 2014. doi:10.1038/ncomms6424.
- [6] Adrian Baddeley, Pablo Gregori, Jorge Mateu, Radu Stoica, and Dietrich Stoyan. *Case Studies in Spatial Point Process Modeling*, volume 185. Springer, 2006. doi:10.1007/0-387-31144-0.
- [7] Bruce J Berne and Philip Pechukas. Gaussian model potentials for molecular interactions. *The Journal of Chemical Physics*, 56(8):4213–4216, 1972.
- [8] Steffen Betsch and Günter Last. On the uniqueness of gibbs distributions with a non-negative and subcritical pair potential. *arXiv preprint arXiv:2108.06303*, 2021.
- [9] Karl Bringmann, Ralph Keusch, and Johannes Lengler. Geometric inhomogeneous random graphs. *Theoretical Computer Science*, 760:35–54, 2019. doi:10.1016/j.tcs.2018.08.014.
- [10] Xiaoyu Chen, Weiming Feng, Yitong Yin, and Xinyuan Zhang. Rapid mixing of glauber dynamics via spectral independence for all degrees. In *2021 IEEE 62nd Annual Symposium on Foundations of Computer Science (FOCS)*, pages 137–148, 2022. doi:10.1109/FOCS52979.2021.00022.
- [11] 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, 2020. doi:10.1109/FOCS46700.2020.00124.
- [12] 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. doi:10.1145/3406325.3451035.
- [13] Daryl J. Daley and David Vere-Jones. *An Introduction to the Theory of Point Processes. Volume II: General Theory and Structure*. Springer, 2008. doi:10.1111/j.1751-5823.2008.00054_18.x.
- [14] Bradley Efron and Charles Stein. The jackknife estimate of variance. *The Annals of Statistics*, pages 586–596, 1981. doi:10.1214/aos/1176345462.
- [15] Fouad El Azhar, Marc Baus, Jean-Paul Ryckaert, and Evert Jan Meijer. Line of triple points for the hardcore yukawa model: A computer simulation study. *The Journal of Chemical Physics*, 112(11):5121–5126, 2000.
- [16] Weiming Feng, Heng Guo, and Yitong Yin. Perfect sampling from spatial mixing. *Random Structures & Algorithms*, 61(4):678–709, 2022.
- [17] Tobias Friedrich, Andreas Göbel, Maximilian Katzmann, Martin S. Krejca, and Marcus Pappik. Algorithms for hard-constraint point processes via discretization. In *Computing and Combinatorics: 28th International Conference, COCOON 2022, Shenzhen, China, October 22–24, 2022, Proceedings*, pages 242–254. Springer, 2023.
- [18] Tobias Friedrich, Andreas Göbel, Martin Krejca, and Marcus Pappik. A spectral independence view on hard spheres via block dynamics. In *48th International Colloquium on Automata, Languages, and Programming (ICALP)*, volume 198, pages 66:1–66:15, 2021. doi:10.4230/LIPIcs.ICALP.2021.66.

- [19] Nancy L. Garcia. Perfect simulation of spatial processes. *Resenhas do Instituto de Matemática e Estatística da Universidade de São Paulo*, 4(3):283–325, 2000.
- [20] Hans-Otto Georgii and Torsten Kühneth. Stochastic comparison of point random fields. *Journal of Applied Probability*, 34(4):868–881, 1997. doi:10.2307/3215003.
- [21] Heng Guo and Mark Jerrum. Perfect simulation of the hard disks model by partial rejection sampling. *Annales de l’Institut Henri Poincaré D*, 8(2):159–177, 2021. doi:10.4171/AIHPD/99.
- [22] Olle Häggström, Marie-Colette NM Van Lieshout, and Jesper Møller. Characterization results and markov chain monte carlo algorithms including exact simulation for some spatial point processes. *Bernoulli*, pages 641–658, 1999. doi:10.2307/3318694.
- [23] Tyler Helmuth, Will Perkins, and Samantha Petti. Correlation decay for hard spheres via markov chains. *The Annals of Applied Probability*, 32(3):2063–2082, 2022.
- [24] Mark Huber. Spatial birth–death swap chains. *Bernoulli*, pages 1031–1041, 2012.
- [25] Sabine Jansen. Gibbsian point processes, 2018. URL: <https://www.math.lmu.de/~jansen/gibbspp.pdf>.
- [26] Matthew Jenssen, Marcus Michelen, and Mohan Ravichandran. Quasipolynomial-time algorithms for gibbs point processes. *Combinatorics, Probability and Computing*, pages 1–15, 2023.
- [27] Dmitri Krioukov, Fragkiskos Papadopoulos, Maksim Kitsak, Amin Vahdat, and Marián Boguná. Hyperbolic geometry of complex networks. *Physical Review E*, 82(3):036106, 2010. doi:10.1103/PhysRevE.82.036106.
- [28] Liang Li, Pinyan Lu, and Yitong Yin. Correlation decay up to uniqueness in spin systems. In *Proceedings of the twenty-fourth annual ACM-SIAM symposium on Discrete algorithms (SODA)*, pages 67–84, 2013. URL: <https://dlnext.acm.org/doi/10.5555/2627817.2627822>.
- [29] László Lovász. *Large networks and graph limits*, volume 60. American Mathematical Soc., 2012.
- [30] Colin McDiarmid. On the method of bounded differences. *Surveys in combinatorics*, 141(1):148–188, 1989. doi:10.1017/CBO9781107359949.008.
- [31] Colin McDiarmid. Concentration. In *Probabilistic methods for algorithmic discrete mathematics*, pages 195–248. Springer, 1998. doi:10.1007/978-3-662-12788-9_6.
- [32] Nicholas Metropolis, Arianna W. Rosenbluth, Marshall N. Rosenbluth, Augusta H. Teller, and Edward Teller. Equation of state calculations by fast computing machines. *The Journal of Chemical Physics*, 21(6):1087–1092, 1953. doi:10.1063/1.1699114.
- [33] Marcus Michelen and Will Perkins. Potential-weighted connective constants and uniqueness of gibbs measures. *CoRR*, abs/2109.01094, 2021. URL: <https://arxiv.org/abs/2109.01094>.
- [34] Marcus Michelen and Will Perkins. Strong spatial mixing for repulsive point processes. *Journal of Statistical Physics*, 189(1):9, 2022.
- [35] Jesper Møller. On the rate of convergence of spatial birth-and-death processes. *Annals of the Institute of Statistical Mathematics*, 41(3):565–581, 1989. doi:10.1007/BF00050669.

- [36] Jesper Møller and Rasmus P Waagepetersen. Modern statistics for spatial point processes. *Scandinavian Journal of Statistics*, 34(4):643–684, 2007. doi:10.1111/j.1467-9469.2007.00569.x.
- [37] Jesper Møller and Rasmus Plenge Waagepetersen. *Statistical inference and simulation for spatial point processes*. CRC press, 2003.
- [38] JS Rowlinson. The yukawa potential. *Physica A: Statistical Mechanics and its Applications*, 156(1):15–34, 1989.
- [39] David Ruelle. *Statistical Mechanics: Rigorous Results*. World Scientific, 1999.
- [40] Alistair Sinclair, Piyush Srivastava, Daniel Štefankovič, and Yitong Yin. Spatial mixing and the connective constant: Optimal bounds. *Probability Theory and Related Fields*, 168(1):153–197, 2017. doi:10.1007/s00440-016-0708-2.
- [41] Alistair Sinclair, Piyush Srivastava, and Marc Thurley. Approximation algorithms for two-state anti-ferromagnetic spin systems on bounded degree graphs. *Journal of Statistical Physics*, 155(4):666–686, 2014. doi:10.1007/s10955-014-0947-5.
- [42] Alistair Sinclair, Piyush Srivastava, and Yitong Yin. Spatial mixing and approximation algorithms for graphs with bounded connective constant. In *2013 IEEE 54th Annual Symposium on Foundations of Computer Science*, pages 300–309. IEEE, 2013.
- [43] 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 (FOCS)*, pages 361–369, 2012. doi:10.1109/FOCS.2012.56.
- [44] Daniel Štefankovič, Santosh Vempala, and Eric Vigoda. Adaptive simulated annealing: A near-optimal connection between sampling and counting. *Journal of the ACM*, 56(3):1–36, 2009. doi:10.1145/1516512.1516520.
- [45] Marie-Colette NM Van Lieshout. *Markov point processes and their applications*. World Scientific, 2000.
- [46] Eric Vigoda. A note on the glauber dynamics for sampling independent sets. *The Electronic Journal of Combinatorics*, 8(1):1–8, 2001. doi:10.37236/1552.
- [47] 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.
- [48] Takeshi Yoshida and Shiro Kamakura. Liquid-solid transitions in systems of soft repulsive forces: Softness of potentials and a maximum in melting curves. *Progress of Theoretical Physics*, 52(3):822–839, 1974. doi:10.1143/PTP.52.822.