A spectral independence view on hard spheres via block dynamics

Tobias Friedrich
Hasso Plattner Institute, University of Potsdam, Potsdam, Germany

Andreas Göbel
Hasso Plattner Institute, University of Potsdam, Potsdam, Germany

Martin S. Krejca
Sorbonne University, CNRS, LIP6, Paris, France

Marcus Pappik
Hasso Plattner Institute, University of Potsdam, Potsdam, Germany

Abstract

The hard-sphere model is one of the most extensively studied models in statistical physics. It describes the continuous distribution of spherical particles, governed by hard-core interactions. An important quantity of this model is the normalizing factor of this distribution, called the *partition function*. We propose a Markov chain Monte Carlo algorithm for approximating the grand-canonical partition function of the hard-sphere model in d dimensions. Up to a fugacity of $\lambda < \mathrm{e}/2^d$, the runtime of our algorithm is polynomial in the volume of the system. This covers the entire known real-valued regime for the uniqueness of the Gibbs measure.

Key to our approach is to define a discretization that closely approximates the partition function of the continuous model. This results in a discrete hard-core instance that is exponential in the size of the initial hard-sphere model. Our approximation bound follows directly from the correlation decay threshold of an infinite regular tree with degree equal to the maximum degree of our discretization. To cope with the exponential blow-up of the discrete instance we use clique dynamics, a Markov chain that was recently introduced in the setting of abstract polymer models. We prove rapid mixing of clique dynamics up to the tree threshold of the univariate hard-core model. This is achieved by relating clique dynamics to block dynamics and adapting the spectral expansion method, which was recently used to bound the mixing time of Glauber dynamics within the same parameter regime.

2012 ACM Subject Classification Theory of computation → Random walks and Markov chains

Keywords and phrases Hard-sphere Model, Markov Chain, Partition Function, Gibbs Distribution, Approximate Counting, Spectral Independence

Digital Object Identifier 10.4230/LIPIcs.ICALP.2021.61

Category Track A: Algorithms, Complexity and Games

Related Version *Full Version*: https://arxiv.org/abs/2102.07443

Funding *Martin S. Krejca*: This work was supported by the Paris Île-de-France Region.

1 Introduction

Statistical physics models particle systems as probability distributions. One of the most fundamental and mathematically challenging models in this area is the hard-sphere model, which plays a central role in understanding the thermodynamic properties of monoatomic gases and liquids [7, 29]. It is a continuous model that studies the distribution and macroscopic behavior of indistinguishable spherical particles, assuming only hard-core interactions, i.e., no two particles can occupy the same space.

We focus on computational properties of the *grand-canonical ensemble* of the hard-sphere model in a finite d-dimensional cubic region $\mathbb{V} = [0, \ell)^d$ in the Euclidean space. In the

grand-canonical ensemble, the system can exchange particles with its surrounding based on a fugacity parameter λ, which is inverse to the temperature of the system. For the rest of the paper, we make the common assumption that the system is normalized such that the particles have unit volume. That means we fix their radii to $r = (1/v_d)^{1/d}$, where v_d is the volume of a unit sphere in d dimensions.

A simple probabilistic interpretation of the distribution of particles in the grand-canonical ensemble is that centers of points that are drawn from a Poisson point process on $\mathbb{V}$ with intensity λ, conditioned on the event that no two particles overlap (i.e., every pair of centers has distance at least $2r$). The resulting distribution over particle configurations in $\mathbb{V}$ is called the *Gibbs distribution* of the model. An important quantity in such models is the so called partition function $Z(\mathbb{V}, \lambda)$, which can be seen as the normalizing constant of the Gibbs distribution. Formally, it is defined as

$$Z(\mathbb{V}, \lambda) = 1 + \sum_{k \in \mathbb{N}_{>0}} \frac{\lambda^k}{k!} \int_{\mathbb{V}^k} D\Big(x^{(1)}, \dots, x^{(k)}\Big) \, \mathrm{d}\nu^{d \times k},$$

where

$$D\Big(x^{(1)}, \dots, x^{(k)}\Big) = \begin{cases} 1 \text{ if } d\big(x^{(i)}, x^{(j)}\big) \geq 2r \text{ for all } i, j \in [k] \text{ with } i \neq j \\ 0 \text{ otherwise} \end{cases}$$

and $\nu^{d \times k}$ is the Lebesgue measure on $\mathbb{R}^{d \times k}$. Commonly, two computational task are associated with the grand-canonical hard-sphere model: (1) approximating its partition function $Z(\mathbb{V}, \lambda)$, and (2) approximately sampling from the Gibbs distribution.

Studying computational aspects of the hard-sphere model carries a historical weight, as in the seminal work of Metropolis [41], the Monte Carlo method was introduced to investigate a two-dimensional hard-sphere model. Approximate-sampling Markov chain approaches have been mainly focused on the canonical ensemble of the model, that is, the system does not exchange particles with its surrounding and thus, the distribution is defined over a fixed number of spheres [31, 36, 34]. Considering the grand canonical ensemble, exact sampling algorithms have appeared in the literature for the two-dimensional model without asymptotic runtime guarantees [37, 38, 43]. A result that is more aligned with theoretical computer science was given in [28], where the authors introduced an exact sampling algorithm for the grand-canonical hard-sphere model in d-dimensions. Their algorithm is based on partial rejection sampling with a runtime linear in the volume of the system $|\mathbb{V}|$ when assuming a continuous computational model and access to a sampler from a continuous Poisson point process. Their approach is guaranteed to apply for $\lambda < 2^{-(d+1/2)}$.

Besides such sampling results, there is an ongoing effort to improve the known fugacity regime where the Gibbs measure is unique and correlations decay exponentially fast [22, 14, 32, 42]. Note that for many discrete spin systems, such as the hard-core model, correlation decay is closely related to the applicability of different methods for efficient approximation of the partition function [50, 24, 54]. Recently, the correlation decay bounds for the hard-sphere model were improved in [32] to $\lambda < 2/2^d$, using probabilistic arguments, and in [42] to $\lambda < \mathrm{e}/2^d$, based on an analytic approach. A common feature of [32] and [42] is that they translated tools originally developed in theoretical computer science for investigating the discrete hard-core model to the continuous domain.

Our work is in line with the computational view on the hard-sphere model but more algorithmic in nature. We investigate the range of the fugacity λ for which an approximation of $Z(\mathbb{V}, \lambda)$ can be obtained efficiently in terms of the volume of the system $|\mathbb{V}|$, assuming a

discrete computational model. Our main result is that for all $\lambda < \mathrm{e}/2^d$ there is a randomized algorithm for ε-approximating the partition function in time polynomial in $|\mathbb{V}|$ and $1/\varepsilon$.

▶ **Theorem 1.** *Let $(\mathbb{V}, \lambda)$ be an instance of the continuous hard-sphere model with $\mathbb{V} = [0, \ell)^d$. If there is a $\delta \in (0, 1]$ such that*

$$\lambda \leq (1 - \delta)\frac{\mathrm{e}}{2^d},$$

then for each $\varepsilon \in (0, 1]$ there is a randomized ε-approximation of $Z(\mathbb{V}, \lambda)$ computable in time polynomial in $|\mathbb{V}|^{1/\delta^2}$ and $\frac{1}{\varepsilon}$.

Note that this bound on λ precisely coincides with the best known bound for the uniqueness of the Gibbs measure in the thermodynamic limit, recently established in [42]. For many discrete spin systems, such as the hard-core model or general anti-ferromagnetic 2-state spin systems, the region of efficient approximation of the partition function is closely related to uniqueness of the Gibbs measure. More precisely, it can be shown that the partition function of every graph of maximum degree Δ can be approximated efficiently if the corresponding Gibbs distribution on an infinite Δ regular tree is unique [39, 53]. A detailed discussion for the discrete hard-core model can be found in the next subsection. In a sense, Theorem 1 can be seen as the algorithmic counterpart of the recent uniqueness result for the continuous hard-sphere model. This answers an open question, asked in [42].

The way we prove our result is quite contrary to [32] and [42]. Instead of translating discrete tools from computer science into the continuous domain, we rather discretize the hard-sphere model. By this, existing algorithmic and probabilistic techniques for discrete models become available, and we avoid continuous analysis.

Our applied discretization scheme is fairly intuitive and results in an instance of the discrete hard-core model. This model has been extensively studied in the computer science community. However, as this hard-core instance is exponential in the size of the continuous system $|\mathbb{V}|$, existing approaches for approximating its partition function, such a Markov chain Monte Carlo methods based on Glauber dynamics, are not feasible. We overcome this problem by applying a Markov chain Monte Carlo approach based on clique dynamics, which were introduced in [23] in the setting of abstract polymer models. Previously known conditions for the rapid mixing of clique dynamics were developed for the multivariate version of the hard-core model. Due to this generality, those conditions do not result in the desired bound in our univariate setting. Instead we relate those clique dynamics to another Markov chain, called *block dynamics*. We then prove the desired mixing time for the block dynamics by adapting a recently introduced technique for bounding the mixing time of Markov chains, based on local spectral expansion [2]. Together with a known self-reducibility scheme for clique dynamics, this results in the desired approximation algorithm.

Note that we aim for a rigorous algorithmic result for approximating the partition function of the continuous hard-sphere model. To be in line with commonly used discrete computational models, our Markov chain Monte Carlo algorithm does not assume access to a continuous sampler but instead samples approximately from a discretized version of the Gibbs distribution. Note that sampling from the continuous hard-sphere partition function cannot be done using a discrete computation model as this would involve infinite float pointer precision. For practical matters, our discretization of the Gibbs distribution can be seen as an approximation of the original continuous Gibbs measure. However, a rigorous comparison between both distributions based on total variation distance is not applicable, due to the fact that one is discrete whereas the other is continuous in nature.

Assuming access to a continuous sampler, we believe that our approach can be used to obtain an approximation of the Gibbs distribution of the continuous model within the same fugacity regime, by adding small perturbations to the drawn sphere centers. This would be in line with the relation between the mixing time of continuous heat-bath dynamics and strong spatial mixing, pointed out in [32], combined with the uniqueness bound from [42].

In Sections 1.1–1.3 we discuss our contributions in more detail and explain how they relate to the existing literature. Finally in Section 1.4 we discuss possible extensions and future work. All technical details, statements and proofs are presented the full version of the paper.

1.1 Discretization and hard-core model

Our discretization scheme expresses the hard-sphere partition function as the partition function of an instance of the (univariate) hard-core model. An instance of the hard-core model is a tuple (G, λ) where $G = (V, E)$ is an undirected graph and $\lambda \in \mathbb{R}_{>0}$. Its partition function is defined as

$$Z(G, \lambda) \coloneqq \sum_{I \in \mathcal{I}(G)} \lambda^{|I|},$$

where $\mathcal{I}(G)$ denotes the independent sets of G. A common way to obtain an approximation for the partition function is by applying a Markov chain Monte Carlo algorithm. This involves sampling from the Gibbs distribution $\mu^{(G,\lambda)}$ of (G, λ), which is a probability distribution on $\mathcal{I}(G)$ that assigns each independent set $I \in \mathcal{I}(G)$ the probability

$$\mu^{(G,\lambda)}(I) = \frac{\lambda^{|I|}}{Z(G, \lambda)}.$$

Conditions for efficient approximation of the hard-core partition function have been studied extensively in the theoretical computer science community. Due to hardness results in [50] and [24], it is known that for general graphs of maximum degree $\Delta \in \{3\} \cup \mathbb{N}_{>5}$ there is a critical parameter value $\lambda_c(\Delta) = (\Delta - 1)^{\Delta-1}/(\Delta - 2)^{\Delta}$, such that there is no FPRAS for the partition function of (G, λ) for $\lambda > \lambda_c(\Delta)$, unless RP = NP. On the other hand, in [54] it was proven that there is a deterministic algorithm for approximating the partition function of (G, λ) for $\lambda < \lambda_c(\Delta)$ that runs in time $|V|^{O(\Delta)}$. The critical value $\lambda_c(\Delta)$ is especially interesting, as it precisely coincides with the upper bound on λ for which the hard-core model on an infinite Δ-regular tree exhibits strong spatial mixing and a unique Gibbs distribution [54]. For this reason, it is also referred to as the *tree threshold*. This relation between computational hardness and phase transition in statistical physics is one of the most celebrated results in the area. Both, the hardness results [25, 3] and the approximation algorithms [46, 30] were later generalized for complex λ.

Note that the computational hardness above the tree threshold $\lambda_c(\Delta)$ for general graphs of maximum degree Δ applies to both, randomized and deterministic algorithms. However, in the randomized setting, Markov chain Monte Carlo methods are known to improve the runtime of the algorithm introduced in [54]. Those approaches use the vertex-wise self-reducibility of the hard-core model to construct a randomized approximation of the partition function based on an approximate sampler for the Gibbs distribution. Commonly, a Markov chain on the state space $\mathcal{I}(G)$, called *Glauber dynamics*, is used to construct the sampling scheme. At each step, a vertex $v \in V$ is chosen uniformly at random. With probability $\lambda/(1 + \lambda)$ the chain tries to add v to the current independent set and otherwise it tries to remove it. The resulting Markov chain is ergodic and reversible with respect

to the Gibbs distribution, meaning that it eventually converges to $\mu^{(G,\lambda)}$. A sequence of results has shown that for all $\Delta \geq 3$ there is a family of graphs with maximum degree Δ, such that the Glauber dynamics are torpidly mixing for $\lambda > \lambda_c(\Delta)$, even without additional complexity-theoretical assumptions [17, 27, 45]. Whether the Glauber dynamics are rapidly mixing for the entire regime $\lambda < \lambda_c(\Delta)$ remained a long-standing open problem, until recently the picture was completed [2]. By relating spectral expansion properties of certain random walks on simplicial complexes to the Glauber dynamics, it was shown that the mixing time is polynomial in $|V|$ below the tree threshold. The mixing time was recently further improved in [12] for a broader class of spin systems by combining simplicial complexes with entropy factorization and using the modified log-Sobolev inequality.

By mapping the hard-sphere model to an instance of the hard-core model we can make use of the existing results about approximation and sampling below the tree threshold. Roughly, our discretization scheme restricts the positions of sphere centers to an integer grid, while scaling the radii of spheres and the fugacity appropriately. For a hard-sphere instance $(\mathbb{V}, \lambda)$ with $\mathbb{V} = [0, \ell)^d$ the hard-core representation for resolution $\rho \in \mathbb{R}_{\geq 1}$ is a hard-core instance (G_ρ, λ_ρ) with $G_\rho = (V_\rho, E_\rho)$. Each vertex $v \in V_\rho$ represents a grid point in the finite integer lattice of side length $\rho\ell$. Two distinct vertices in V_ρ are connected by an edge in E_ρ if the Euclidean distance of the corresponding grid points is less than $2\rho r$. Furthermore, we set $\lambda_\rho = \lambda/\rho^d$. We provide the following result on the rate of convergence of $Z(G_\rho, \lambda_\rho)$ to the hard-sphere partition function $Z(\mathbb{V}, \lambda)$ in terms of ρ.

▶ **Lemma 2.** *Let $(\mathbb{V}, \lambda)$ be an instance of the continuous hard-sphere model in d dimensions. For each resolution $\rho \geq 2\sqrt{d}$ it holds that*

$$1 - \rho^{-1}\mathrm{e}^{\Theta(|\mathbb{V}| \ln|\mathbb{V}|)} \leq \frac{Z(\mathbb{V}, \lambda)}{Z(G_\rho, \lambda_\rho)} \leq 1 + \rho^{-1}\mathrm{e}^{\Theta(|\mathbb{V}| \ln|\mathbb{V}|)}.$$

Although proving this rate of convergence involves some detailed geometric arguments, there is an intuitive reason why the partition functions converge eventually as $\rho \to \infty$. Increasing the resolution ρ also linearly increases the side length of the grid and the minimum distance that sphere centers can have. This is equivalent to putting a grid into $\mathbb{V}$ with increasing granularity but fixing the radii of spheres instead. However, only scaling the granularity of this grid increases the number of possible configurations by roughly ρ^d, which would cause the partition function of the hard-core model to diverge as $\rho \to \infty$. To compensate for this, we scale the weight of each vertex in the hard-core model by the inverse of this factor.

Using this discretization approach, the fugacity bound from Theorem 1 results from simply considering Δ_ρ, the maximum degree of G_ρ and comparing λ_ρ with the corresponding tree threshold $\lambda_c(\Delta_\rho)$. Recall that we assume $r = (1/v_d)^{1/d}$. A simple geometric argument shows that Δ_ρ is roughly upper bounded by $2^d\rho^d$ for sufficiently large ρ. Now, observe that

$$\lambda_\rho = \frac{\lambda}{\rho^d} < \lambda_c\big(2^d\rho^d\big),$$

for $\lambda < \rho^d\lambda_c\big(2^d\rho^d\big)$. This follows from the fact that $\rho^d\lambda_c\big(2^d\rho^d\big)$ converges to $\mathrm{e}/2^d$ from above for $\rho \to \infty$. Thus, the approximation bound from Theorem 1 and the uniqueness bound in [42] coincide with the regime of λ, for which λ_ρ is below the tree threshold $\lambda_c(\Delta_\rho)$ in the limit $\rho \to \infty$.

The arguments above show that for a sufficiently high resolution ρ the partition function of the hard-sphere model $Z(\mathbb{V}, \lambda)$ is well approximated by the partition function of our discretization (G_ρ, λ_ρ) and that (G_ρ, λ_ρ) is below the tree threshold for $\lambda < \mathrm{e}/2^d$. However,

this does not immediately imply an approximation algorithm within the desired runtime bounds. Based on Lemma 2, we still need to choose ρ exponentially large in the volume $|\mathbb{V}|$. Note that the number of vertices in G_ρ is roughly $|V_\rho| \in \Theta\big(\rho^d|\mathbb{V}|\big)$. Even without explicitly constructing the graph, this causes problems, as the best bound for the mixing time of the Glauber dynamics is polynomial in $|V_\rho|$ and thus exponential in $|\mathbb{V}|$. Intuitively, the reason for this mixing time is that the Glauber dynamics only change one vertex at each step. Assuming that each vertex should be updated at least once to remove correlations with the initial state, any mixing time that is sublinear in the number of vertices is unlikely. We circumvent this problem by applying dynamics that update multiple vertices at each step but still allow each step to be computed efficiently without constructing the graph explicitly.

1.2 Block and clique dynamics

Most of the results that we discuss from now on apply to the multivariate version of the hard-core model, that is, each vertex $v \in V$ has its own weight λ_v. For a given graph $G = (V, E)$ we denote the set of such vertex weights by $\boldsymbol{\lambda} = \{\lambda_v\}_{v\in V}$ and write $(G, \boldsymbol{\lambda})$ for the resulting multivariate hard-core instance. In the multivariate setting, the contribution of an independent set $I \in \mathcal{I}(G)$ to the partition function is defined as the product of its vertex weights (i.e., $\prod_{v\in I} \lambda_v$), where the contribution of the empty set is fixed to 1. Similar to the univariate hard-core model, the Gibbs distribution assigns a probability to each independent set proportionally to its contribution to the partition function.

As we discussed before, the main problem with approximating the partition function of our discretization (G_ρ, λ_ρ) is that the required graph G_ρ is exponential in the volume of the original continuous system $|\mathbb{V}|$. As the Glauber dynamics Markov chain only updates a single vertex at each step, the resulting mixing time is usually polynomial in the size of the graph, which is not feasible in our case. Various extensions to Glauber dynamics for updating multiple vertices in each step have been proposed in the literature, two of which we discuss in the following.

Clique dynamics

Recently, in [23] a Markov chain, called *clique dynamics*, was introduced in order to efficiently sample from the Gibbs distribution of abstract polymer models. Note that this is similar to our algorithmic problem, as abstract polymer models resemble multivariate hard-core instances. For a given graph $G = (V, E)$, we call a set $\Lambda = \{\Lambda_i\}_{i\in[m]} \subseteq 2^V$ a clique cover of size m if and only if its union covers all vertices V and each $\Lambda_i \in \Lambda$ induces a clique in G. For an instance of the multivariate hard-core model $(G, \boldsymbol{\lambda})$ and a given clique cover Λ of G with size m the clique dynamics Markov chain $\mathcal{C}(G, \boldsymbol{\lambda}, \Lambda)$ is defined as follows. First, a clique $\Lambda_i \in \Lambda$ for $i \in [m]$ is chosen uniformly at random. Let us write $G[\Lambda_i]$ for the subgraph, induced by Λ_i, and $\boldsymbol{\lambda}[\Lambda_i] = \{\lambda_v\}_{v\in\Lambda_i}$ for the corresponding set of vertex weights. Next, an independent set from $\mathcal{I}(G[\Lambda_i])$ is chosen according to the Gibbs distribution $\mu^{(G[\Lambda_i],\boldsymbol{\lambda}[\Lambda_i])}$. Note that, as the vertices Λ_i form a clique, such an independent set is either the empty set or contains a single vertex from $v \in \Lambda_i$. If the empty set is drawn, all vertices from Λ_i are removed from the current independent set. Otherwise, if a single vertex $v \in \Lambda_i$ is drawn, the chain tries to add v to the current independent set.

Using a coupling argument, it was proven in [23] that the so-called *clique dynamics condition* implies that for any clique cover of size m the clique dynamics are mixing in time polynomial in m and $Z_{\max}$, where $Z_{\max} = \max_{i\in[m]}\{Z(G[\Lambda_i], \boldsymbol{\lambda}[\Lambda_i])\}$ denotes the maximum partition function of a clique in Λ. This is important for the application to polymer models, as

they are usually used to model partition functions of other spin systems, which often results in a multivariate hard-core model of exponential size [33, 9, 35, 10, 6, 8, 26]. As discussed in [23], those instances tend to have polynomial size clique covers that arise naturally from the original spin system. In such cases, the mixing time of clique dynamics is still polynomial in the size of original spin system, as long as the clique dynamics condition is satisfied.

This is very similar to our discretization (G_ρ, λ_ρ). To see this, set $a = \frac{2\rho}{\sqrt{d}}r$ and divide the d-dimensional integer lattice of side length $\rho\ell$ into cubic regions of side length a. Every pair of integer points within such a cubic region has Euclidean distance less than $2\rho r$, meaning that the corresponding vertices in G_ρ are adjacent. Thus, each such cubic region forms a clique, resulting in a clique cover of size $(\rho\ell/a)^d \in \mathrm{O}(|\mathbb{V}|)$. This means, there is always a clique cover with size linear in $|\mathbb{V}|$ and independent of the resolution ρ. By showing that, for the univariate hard-core model, the mixing time of clique dynamics is polynomial in the size of the clique cover for all $\lambda_\rho < \lambda_c(\Delta_\rho)$, we obtain a Markov chain with mixing time polynomial in $|\mathbb{V}|$ independent of the resolution ρ. Unfortunately, the clique dynamics condition does not hold for the entire regime up to the tree threshold in the univariate hard-core model. We overcome this problem by proving a new condition for rapid mixing of clique dynamics based on a comparison with block dynamics.

Block dynamics

Block dynamics are a very natural generalization of Glauber dynamics to arbitrary sets of vertices. For a given graph $G = (V, E)$, we call a set $\Lambda = \{\Lambda_i\}_{i\in[m]} \subseteq 2^V$ a block cover of size m if and only if its union covers all vertices V. We refer to the elements of Λ as blocks. Note that the clique cover discussed before is a special case of a block cover, where all blocks are restricted to be cliques. At each step, the block dynamics Markov chain $\mathcal{B}(G, \lambda, \Lambda)$ chooses a block $\Lambda_i \in \Lambda$ uniformly at random. Then, the current independent set is updated on Λ_i based on the projection of the Gibbs distribution onto Λ_i and conditioned on the current independent set outside Λ_i.

In fact, block dynamics are defined for a much more general class of spin systems than the hard-core model. However, due to the fact that each step of the Markov chain involves sampling from a conditional Gibbs distribution, block dynamics are rarely used as an algorithmic tool on its own. Instead, they are usually used to deduce rapid mixing of other dynamics.

For spin systems on lattice graphs, close connections between the mixing time of block dynamics and Glauber dynamics are known [40]. Such connections were for example applied to improve the mixing time of Glauber dynamics of the Monomer Dimer model on torus graphs [51]. Moreover, block dynamics were used to improve conditions for rapid mixing of Glauber dynamics on specific graph classes, such as proper colorings [16, 18, 19, 44] or the hard-core model [18, 44] in sparse random graphs. A very general result for the mixing time of block dynamics was achieved in [4], who proved that for all spin systems on a finite subgraph of the d-dimensional integer lattice the mixing time of block dynamics is polynomial in the number of blocks if the spin system exhibits strong spatial mixing. This result was later generalized in [5] for the Ising model on arbitrary graphs. Very recently, block dynamics based random equally-sized blocks where used in [12] to prove entropy factorization and improve the mixing time of Glauber dynamics for a variety of discrete spin systems up to the tree threshold.

Although our discretization works by restricting sphere positions to the integer lattice, the resulting graph is rather different from the lattice. Thus, results like those in [4] do not apply to our setting. However, on the other hand, we do not need to prove rapid mixing for

arbitrary block covers. Instead, in order to obtain rapid mixing for clique dynamics, it is sufficient to establish this result for cases where all blocks are cliques.

Applying block dynamics directly would involve sampling from a conditional Gibbs distribution within each clique. Due to the exponential size of the cliques in our discretization, this would impose additional algorithmic challenges. Instead, similar to the previous literature, we rather use block dynamics as a tool for proving rapid mixing of another Markov chain, namely clique dynamics.

Improved mixing condition for clique dynamics via block dynamics

We analyze the mixing time of clique dynamics for a given clique cover by relating it to the mixing time of block dynamics, using the cliques as blocks. This is done by investigating a notion of pairwise influence between vertices that has also been used to establish rapid mixing of Glauber dynamics up to the tree threshold [2]. Let $\mathbb{P}_G[w]$ denote the probability of the event that a vertex $w \in V$ is in an independent set drawn from $\mu^{(G,\boldsymbol{\lambda})}$. Further, let $\mathbb{P}_G[\overline{w}]$ denote the probability that w is not in an independent set. We extend this abuse of notation to conditional probabilities, so $\mathbb{P}_G[\cdot \mid \overline{w}]$ for example denotes the probability of some event conditioned on w not being in an independent set. For a pair of vertices $v, w \in V$ the influence $\Psi_G(v, w)$ of v on w is defined as

$$\Psi_G(v,w) = \begin{cases} 0 & \text{if } v = w, \\ \mathbb{P}_G[w \mid v] - \mathbb{P}_G[w \mid \overline{v}] & \text{otherwise.} \end{cases}$$

The following condition in terms of pairwise influence is central to our analysis.

▶ **Condition 3.** *Let $(G, \boldsymbol{\lambda})$ be an instance of the multivariate hard-core model. There is a constant $C \in \mathbb{R}_{>0}$ and a function $q\colon V \to \mathbb{R}_{>0}$ such that for all $S \subseteq V$ and $r \in S$ it holds that*

$$\sum_{v \in S} |\Psi_G(r,v)| q(v) \leq C q(r).$$

Note that this condition appeared before in [13], where it was used for bounding the mixing time of Glauber dynamics for anti-ferromagnetic spin systems. Given Condition 3, we obtain the following result for the mixing time of block dynamics based on a clique cover.

▶ **Theorem 4.** *Let $(G, \boldsymbol{\lambda})$ be an instance of the multivariate hard-core model that satisfies Condition 3. Let Λ be a clique cover for G of size m, and let $Z_{\max} = \max_{i \in [m]} \{Z(G[\Lambda_i], \boldsymbol{\lambda}[\Lambda_i])\}$. The mixing time of the block dynamics $\mathcal{B}(G, \boldsymbol{\lambda}, \Lambda)$, starting from $\emptyset \in \mathcal{I}(G)$, is bounded by*

$$\tau_{\mathcal{B}}^{(\emptyset)}(\varepsilon) \leq m^{\mathrm{O}((2+C)C)} Z_{\max}^{\mathrm{O}((2+C)C)} \ln\left(\frac{1}{\varepsilon}\right).$$

Using a bound for the sum of absolute pairwise influences that was recently established in [13], it follows that the univariate hard-core model satisfies Condition 3 up to the tree threshold. As a result, we know that the mixing time of block dynamics is polynomial in m and $Z_{\max}$ for any clique cover of size m. To the best of our knowledge, this is the first result for the mixing time of block dynamics for the univariate hard-core model on general graphs that holds in this parameter range.

As we aim to apply clique dynamics to avoid sampling from the conditional Gibbs distribution in each step, we still need to prove that Theorem 4 also holds in terms of clique dynamics. To this end, we apply a Markov chain comparison argument from [15] to prove

that using clique dynamics instead of block dynamics for the same clique cover Λ increases the mixing time by at most a factor $2Z_{\max}$. The following corollary, which is central for proving Theorem 1, follows immediately.

▶ **Corollary 5.** *Let (G, λ) be an instance of the univariate hard-core model such that the degree of G is bounded by Δ. Let Λ be a given clique cover of size m with $Z_{\max} = \max_{i\in[m]}\{Z(G[\Lambda_i], \lambda)\}$. Denote by $\mathcal{C} = \mathcal{C}(G, \lambda, \Lambda)$ the corresponding clique dynamics. If there is some $\delta \in \mathbb{R}_{>0}$ such that $\lambda \leq (1-\delta)\lambda_c(\Delta)$ then the mixing time of the clique dynamics $\mathcal{C}$, starting from $\emptyset \in \mathcal{I}(G)$, is bounded by*

$$\tau_{\mathcal{C}}^{(\emptyset)}(\varepsilon) \leq m^{\mathrm{O}\left(1/\delta^2\right)} Z_{\max}^{\mathrm{O}\left(1/\delta^2\right)} \ln\left(\frac{1}{\varepsilon}\right).$$

A side journey: comparison to multivariate conditions

In fact, Corollary 5 is sufficient for our application to the hard-sphere model. However, we also aim to set Condition 3 in the context of other conditions for rapid mixing of clique dynamics for the multivariate hard-core model. Note that such a rapid mixing result for clique dynamics caries over to Glauber dynamics by taking each vertex as a separate clique of size 1.

To this end, we compare Condition 3 to a strict version of the clique dynamics condition, originally introduced in [23] in the setting of clique dynamics for abstract polymer models. It turns out that this strict version of the clique dynamics condition directly implies Condition 3. This is especially interesting, as the clique dynamics condition was initially introduced as a local condition (only considering the neighborhood of each vertex) and is based on a coupling argument. However, we show that it can as well be understood as a sufficient condition for the global decay of pairwise influence with increasing distance between vertices.

Formally, we say that the strict clique dynamics condition is satisfied for an instance of the multivariate hard-core model $(G, \boldsymbol{\lambda})$ if there is a function $f\colon V \to \mathbb{R}_{>0}$ and a constant $\alpha \in (0, 1)$ such that for all $v \in V$ it holds that

$$\sum_{w\in N(v)} \frac{\lambda_w}{1+\lambda_w} f(w) \leq (1-\alpha) f(v),$$

where $N(v)$ is the neighborhood of v in G. This is a strict version of the clique dynamics condition in that the original clique dynamics condition would correspond to the case $\alpha = 0$ (i.e., the strict clique dynamics condition requires some strictly positive slack α).

The result of our comparison is summarized in the following statement.

▶ **Lemma 6.** *Let $(G, \boldsymbol{\lambda})$ be an instance of the multivariate hard-core model. If $(G, \boldsymbol{\lambda})$ satisfies the strict clique dynamics condition for a function f and a constant α, then it also satisfies Condition 3 for $q = f$ and $C = \frac{1}{\alpha}$.*

Lemma 6 is proven by translating the calculation of pairwise influences to the self-avoiding walk tree of the graph, based on a result in [13], and applying a recursive argument on this tree.

Despite being an interesting relationship between local coupling arguments and global pairwise influence, Lemma 6 also implies that, from an algorithmic perspective, Theorem 4 can be used to produce similar results as those obtained in [23] for abstract polymer models. Further, note that for the univariate model, using pairwise influence yields strictly better results than any coupling approach in the literature. This raises the question if a refined

argument based on pairwise influences can be used in the multivariate setting to improve on the clique dynamics condition, leading to better approximation results on abstract polymer models.

1.3 Analyzing spectral expansion

As core technique for obtaining Theorem 4, we adapt an approach for bounding the mixing time that was recently used to prove rapid mixing of Glauber dynamics for the entire regime below the tree threshold for several applications, such as the hard-core model [2], general two-state spin systems [13], and proper colorings [11, 21]. The idea is to map the desired distribution to a weighted simplicial complex.

A simplicial complex X over a groundset U is a set family $X \subseteq 2^U$ such that for each $\tau \in X$ every subset of τ is also in X. We call the elements $\tau \in X$ the faces of X and refer to its cardinality $|\tau|$ as dimensionality.

For a univariate hard-core instance (G, λ), the authors of [2] construct a simplicial complex over a ground set U that contains two elements $x_v, x_{\overline{v}} \in U$ for each vertex $v \in V$. For every independent set $I \in \mathcal{I}(G)$, a face $\tau_I \in X$ is introduced such that $x_v \in \tau_I$ if $v \in I$ and $x_{\overline{v}} \in \tau_I$ otherwise. The simplicial complex is completed by taking the downward closure of these faces. Note that by construction all maximum faces of the resulting complex are $|V|$-dimensional and there is a one-to-one correspondence between the maximum faces and the independent sets in $\mathcal{I}(G)$. By assigning each maximum face $\tau_I \in X$ an appropriate weight, the Glauber dynamics can be represented as a random walk on those maximum faces, which is sometimes referred to as the two-step walk or down-up walk. Using a local-to-global theorem [1], the mixing time of this two-step walk can then be bounded based on certain local expansion properties of the simplicial complex X. It is then proved that such local expansion properties are well captured by the largest eigenvalue of the pairwise influence matrix Ψ_G, which is a $|V| \times |V|$ matrix that contains the pairwise influence $\Psi_G(v, w)$ for all $v, w \in V$. Finally, by bounding those influences a bound on this largest eigenvalue of Ψ_G is obtained. This analysis was later refined and generalized in [13] to general two-state spin systems.

This method was independently extended in [11] and [21] to the non-Boolean domain by applying it to the Glauber dynamics for proper colorings. The main differences to the Boolean domain are that elements of the simplicial complex now represent combinations of a vertex and a color. Furthermore, the bound on the local spectral expansion was obtained by using a different influence matrix, which captures the effect of selecting a certain color for one vertex on the distribution of colors for another vertex.

Although we are dealing with the hard-core model, which is Boolean in nature, the way that we model block dynamics is mainly inspired by the existing work on proper colorings [11]. Assume we have an instance of the multivariate hard-core model $(G, \boldsymbol{\lambda})$ and let Λ be a clique cover for G of size m such that every pair of distinct cliques is vertex-disjoint (i.e., Λ is a partition of G into cliques). We construct a simplicial complex X based on a ground set U that contains one element $x_v \in U$ for each vertex $v \in V$ and one additional element $\emptyset_i$ for each clique $\Lambda_i \in \Lambda$. We introduce a face $\tau_I \in X$ for each independent set $I \in \mathcal{I}(G)$ such that for every $\Lambda_i \in \Lambda$ we have $\emptyset_i \in \tau_I$ if $\Lambda_i \cap I = \emptyset$ and $x_v \in \tau_I$ if $\Lambda_i \cap I = \{v\}$ for some $v \in \Lambda_i$. The simplicial complex is completed by taking the downward closure of these faces. All maximum faces of the resulting complex are m-dimensional and there is a bijection between the maximum faces and the independent sets of G. Furthermore, there is a natural partitioning $\{U_i\}_{i \in [m]}$ of the ground set U, each partition U_i corresponding to a clique Λ_i, such that every maximum face in X contains exactly one element from each partition U_i.

By weighting each maximum face of X by the contribution of the corresponding inde-

pendent set to the partition function, the block dynamics based on Λ are equivalent to the two-step walk on X. Thus, in order to bound the mixing time of the block dynamics, it is sufficient to study the local expansion properties of X. To this end, we adapt the influence matrix used for proper colorings in [11]. For $x \in U$, let $\mathbb{P}_G[x]$ denote the probability that $x \in \tau_I$ for an independent set $I \in \mathcal{I}(G)$ drawn from $\mu^{(G,\boldsymbol{\lambda})}$ and corresponding maximum face $\tau_I \in X$. Similarly as for defining pairwise influences, we extend this notation to conditional probabilities. The clique influence matrix $\Phi_{G,\Lambda}$ contains an entry $\Phi_{G,\Lambda}(x,y)$ for each $x, y \in U$ with

$$\Phi_{G,\Lambda}(x,y) = \begin{cases} 0 & \text{if } x,y \in U_i \text{ for some } i \in [m], \\ \mathbb{P}_G[y \mid x] - \mathbb{P}_G[y] & \text{otherwise.} \end{cases}$$

By using similar linear-algebraic arguments as in [11] we prove that the maximum eigenvalue of $\Phi_{G,\Lambda}$ can be used to upper bound the local spectral expansion of X. To obtain Theorem 4 it is then sufficient to relate Condition 3 to the maximum eigenvalue of $\Phi_{G,\Lambda}$. The following lemma establishes this connection.

▶ **Lemma 7.** *Let $(G, \boldsymbol{\lambda})$ be an instance of the multivariate hard-core model that satisfies Condition 3 for a function q and a constant C. For every $S \subseteq V$ and every disjoint clique cover Λ of $G[S]$ it holds that the largest eigenvalue of $\Phi_{G[S],\Lambda}$ is at most $(2 + C)C$.*

Note that our simplicial-complex representation is only given under the assumption that the cliques in the clique cover Λ are pairwise disjoint. Indeed, this is a necessary requirement to map the block dynamics to the two-step walk such that the local-global-theorem from [1] can be applied. Thus, Lemma 7 only helps to prove Theorem 4 for disjoint clique covers. However, we relax this requirement by proving that for every clique cover Λ a disjoint clique cover K can be constructed such that the block dynamics $\mathcal{B}(G, \boldsymbol{\lambda}, \Lambda)$ and $\mathcal{B}(G, \boldsymbol{\lambda}, K)$ have asymptotically the same mixing time. By this comparison argument, we extend Theorem 4 to arbitrary clique covers.

We are aware that, in the case of Glauber dynamics, more recent techniques for combining simplical complex representations with entropy factorization as proposed in [12] yield superior mixing time results. However, in case of the hard-core model, this approach comes with an additional multiplicative factor of $\Delta^{\mathrm{O}(\Delta^2)}$ in the mixing time (see section 8 of [12]). Although negligible for bounded degree graphs, this would be too much for our application, as the degree of our discretization gets exponentially large in the continuous volume $|\mathbb{V}|$ of the system. Thus, directly relating local spectral expansion with the spectral gap of block dynamics is more suitable in our case. We leave as an open question, whether a modification of the approach in [12] can be applied to further improve our mixing time result.

1.4 Outlook

We obtain the fugacity bound from Theorem 1 based on the tree threshold $\lambda_c(\Delta)$ of the hard-core model. An obvious question is whether there are any structural properties of our discretization that can be used to improve this bound. Similar results are known for specific graph classes, such as the 2-dimensional square lattice [48, 52, 54]. In [42] the authors discuss that a generalization of the connective constant to the continuous Euclidean space might be applicable to improve their uniqueness result for the hard-sphere model. A comparable algorithmic result was already established for the discrete hard-core model in [49]. However, any such improvement for our discretization would require the connective constant of G_ρ to be at least by a constant factor small than its maximum degree Δ_ρ. Unfortunately, due to a

result in [47], this is not the case. Although this does not necessarily imply that a similar concept does not work in continuous space, it gives a strong evidence that a more specialized tool instead of the connective constant might be required.

A different direction for future work is to see which other quantities and properties of the model are preserved under discretization. This would especially include the thermodynamic pressure and its analyticity. As a matter of fact, non-analytic points of the pressure along the positive real axis of fugacity in the thermodynamic limit are known to mark phase transitions in infinite volume systems (see for example [42]). One way to approach this could be to prove a relation between zero-freeness of the continuous and the discretized partition function in a complex neighborhood of the real axis by extending our convergence result to the complex domain. Along this line, insights could be gained in how far properties like correlation decay and phase transitions (or their absence) are preserved under sufficiently fine discretization.

From a purely technical point of view, it is interesting to see if our result on the mixing time of block dynamics in Theorem 4 also holds without the requirement of using cliques as blocks. Especially: is the mixing time for block dynamics for the univariate hard-core model polynomial in the number of blocks for any block cover? Most of our techniques that we use for clique covers, such as modeling the distribution as a simplicial complex and relating its local spectral expansion to the clique influence matrix, can be generalized in a straightforward way for different choices of blocks. However, the main difficulty is to relate generalized versions of the clique influence matrix to pairwise influences, as we do in Lemma 7. One way to circumvent this might be to not rely on pairwise influences at all but to rather investigate the influence matrix directly, for example, via different computational-tree methods.

Finally, it would be interesting to see if approaches like ours can be extended to other continuous models from statistical physics (see for example coarse-graining [20]). We believe that the variety of tools that are already established for discrete spin systems are useful in this setting to establish rigorous computational results for different continuous models. We emphasize that clique and block dynamics are a useful computational tool to handle the exponential blow-ups caused by discretization.

References

1 Vedat Levi Alev and Lap Chi Lau. Improved analysis of higher order random walks and applications. In *Proc. of STOC'20*, pages 1198–1211, 2020. doi:10.1145/3357713.3384317.

2 Nima Anari, Kuikui Liu, and Shayan Oveis Gharan. Spectral independence in high-dimensional expanders and applications to the hardcore model. *CoRR*, abs/2001.00303, 2020. URL: http://arxiv.org/abs/2001.00303.

3 Ivona Bezáková, Andreas Galanis, Leslie Ann Goldberg, and Daniel Štefankovič. Inapproximability of the independent set polynomial in the complex plane. In *Proc. of STOC'18*, page 1234–1240, 2018. doi:10.1145/3188745.3188788.

4 Antonio Blanca, Pietro Caputo, Alistair Sinclair, and Eric Vigoda. Spatial mixing and nonlocal Markov chains. *Random Structures & Algorithms*, 55(3):584–614, 2019. doi:10.1002/rsa.20844.

5 Antonio Blanca, Zongchen Chen, and Eric Vigoda. Swendsen-Wang dynamics for general graphs in the tree uniqueness region. *Random Structures & Algorithms*, 56(2):373–400, 2020. doi:10.1002/rsa.20858.

6 Christian Borgs, Jennifer T. Chayes, Tyler Helmuth, Will Perkins, and Prasad Tetali. Efficient sampling and counting algorithms for the Potts model on $\mathbb{Z}^d$ at all temperatures. In *Proc. of STOC'20*, pages 738–751, 2020. doi:10.1145/3357713.3384271.

7 Tomáš Boublik, Ivo Nezbeda, and Karel Hlavaty. *Statistical thermodynamics of simple liquids and their mixtures*. Fundamental Studies in Engineering. Elsevier, 1980.

8 Sarah Cannon and Will Perkins. Counting independent sets in unbalanced bipartite graphs. In *Proc. of SODA'20*, pages 1456–1466, 2020. doi:10.1137/1.9781611975994.88.

9 Katrin Casel, Philipp Fischbeck, Tobias Friedrich, Andreas Göbel, and J. A. Gregor Lagodzinski. Zeros and approximations of holant polynomials on the complex plane. *CoRR*, abs/1905.03194, 2019. URL: http://arxiv.org/abs/1905.03194.

10 Zongchen Chen, Andreas Galanis, Leslie Ann Goldberg, Will Perkins, James Stewart, and Eric Vigoda. Fast algorithms at low temperatures via Markov chains. In *Proc. of APPROX-/RANDOM'19*, pages 41:1–41:14, 2019. doi:10.4230/LIPIcs.APPROX-RANDOM.2019.41.

11 Zongchen Chen, Andreas Galanis, Daniel Stefankovic, and Eric Vigoda. Rapid mixing for colorings via spectral independence. *CoRR*, abs/2007.08058, 2020. URL: http://arxiv.org/abs/2007.08058.

12 Zongchen Chen, Kuikui Liu, and Eric Vigoda. Optimal mixing of Glauber dynamics: Entropy factorization via high-dimensional expansion. *CoRR*, abs/2011.02075, 2020. URL: https://arxiv.org/abs/2011.02075.

13 Zongchen Chen, Kuikui Liu, and Eric Vigoda. Rapid mixing of Glauber dynamics up to uniqueness via contraction. *CoRR*, abs/2004.09083, 2020. URL: http://arxiv.org/abs/2004.09083.

14 Hofer-Temmel Christoph et al. Disagreement percolation for the hard-sphere model. *Electronic Journal of Probability*, 24, 2019.

15 Persi Diaconis and Laurent Saloff-Coste. Comparison theorems for reversible Markov chains. *The Annals of Applied Probability*, pages 696–730, 1993.

16 Martin Dyer, Abraham D. Flaxman, Alan M. Frieze, and Eric Vigoda. Randomly coloring sparse random graphs with fewer colors than the maximum degree. *Random Structures & Algorithms*, 29(4):450–465, 2006. doi:10.1002/rsa.20129.

17 Martin Dyer, Alan Frieze, and Mark Jerrum. On counting independent sets in sparse graphs. *SIAM Journal on Computing*, 31(5):1527–1541, 2002. doi:10.1137/S0097539701383844.

18 Charilaos Efthymiou. MCMC sampling colourings and independent sets of $g(n, d/n)$ near uniqueness threshold. In *Proc. of SODA'14*, page 305–316, 2014. doi:10.1137/1.9781611973402.22.

19 Charilaos Efthymiou, Thomas P. Hayes, Daniel Štefankovič, and Eric Vigoda. Sampling random colorings of sparse random graphs. In *Proc. of SODA'18*, page 1759–1771, 2018. doi:10.1137/1.9781611975031.115.

20 Pep Espanol. Statistical mechanics of coarse-graining. In *Novel Methods in Soft Matter Simulations*, pages 69–115. Springer, 2004. doi:10.1007/978-3-540-39895-0_3.

21 Weiming Feng, Heng Guo, Yitong Yin, and Chihao Zhang. Rapid mixing from spectral independence beyond the Boolean domain. *CoRR*, abs/2007.08091, 2020. URL: http://arxiv.org/abs/2007.08091.

22 Roberto Fernández, Aldo Procacci, and Benedetto Scoppola. The analyticity region of the hard sphere gas. Improved bounds. *Journal of Statistical Physics*, 128(5):1139–1143, 2007. doi:10.1007/s10955-007-9352-7.

23 Tobias Friedrich, Andreas Göbel, Martin S. Krejca, and Marcus Pappik. Polymer dynamics via cliques: New conditions for approximations. *CoRR*, abs/2007.08293, 2020. URL: http://arxiv.org/abs/2007.08293.

24 Andreas Galanis, Qi Ge, Daniel Stefankovic, Eric Vigoda, and Linji Yang. Improved inapproximability results for counting independent sets in the hard-core model. *Random Structures & Algorithms*, 45(1):78–110, 2014. doi:10.1002/rsa.20479.

25 Andreas Galanis, Leslie Ann Goldberg, and Daniel Stefankovic. Inapproximability of the independent set polynomial below the Shearer threshold. *CoRR*, abs/1612.05832, 2016. URL: http://arxiv.org/abs/1612.05832.

26 Andreas Galanis, Leslie Ann Goldberg, and James Stewart. Fast algorithms for general spin systems on bipartite expanders. In *Proc. of MFCS'20*, 2020. To appear. URL: https://arxiv.org/abs/2004.13442.

27 David Galvin and Prasad Tetali. Slow mixing of Glauber dynamics for the hard-core model on regular bipartite graphs. *Random Structures & Algorithms*, 28(4):427–443, 2006. doi:10.1002/rsa.20094.

28 Heng Guo and Mark Jerrum. Perfect simulation of the hard disks model by partial rejection sampling. *CoRR*, abs/1801.07342, 2018. URL: https://arxiv.org/abs/1801.07342v2.

29 Jean-Pierre Hansen and Ian R. McDonald. *Theory of Simple Liquids*. Academic Press, fourth edition edition, 2013. doi:10.1016/B978-0-12-387032-2.00013-1.

30 Nicholas J.A. Harvey, Piyush Srivastava, and Jan Vondrák. Computing the independence polynomial: From the tree threshold down to the roots. In *Proc. of SODA'18*, pages 1557–1576, 2018. doi:10.1137/1.9781611975031.102.

31 Thomas P. Hayes and Cristopher Moore. Lower bounds on the critical density in the hard disk model via optimized metrics. *CoRR*, abs/1407.1930, 2014. URL: http://arxiv.org/abs/1407.1930.

32 Tyler Helmuth, Will Perkins, and Samantha Petti. Correlation decay for hard spheres via Markov chains. *CoRR*, abs/2001.05323, 2020. URL: https://arxiv.org/abs/2001.05323.

33 Tyler Helmuth, Will Perkins, and Guus Regts. Algorithmic Pirogov–Sinai theory. In *Proc. of STOC'19*, pages 1009–1020, 2019. doi:10.1145/3313276.3316305.

34 Matthew Jenssen, Felix Joos, and Will Perkins. On the hard sphere model and sphere packings in high dimensions. *Forum of Mathematics, Sigma*, 7:e1, 2019. doi:10.1017/fms.2018.25.

35 Matthew Jenssen, Peter Keevash, and Will Perkins. Algorithms for #bis-hard problems on expander graphs. In *Proc. of SODA'19*, pages 2235–2247, 2019. doi:10.1137/1.9781611975482.135.

36 Ravi Kannan, Michael W. Mahoney, and Ravi Montenegro. Rapid mixing of several Markov chains for a hard-core model. In *Proc. of ISAAC'03*, pages 663–675, 2003. doi:10.1007/978-3-540-24587-2_68.

37 Wilfrid S. Kendall. Perfect simulation for the area-interaction point process. In *Probability Towards 2000*, pages 218–234. Springer, 1998. doi:10.1007/978-1-4612-2224-8_13.

38 Wilfrid S. Kendall and Jesper Møller. Perfect simulation using dominating processes on ordered spaces, with application to locally stable point processes. *Advances in Applied Probability*, 32(3):844–865, 2000. doi:10.1239/aap/1013540247.

39 Liang Li, Pinyan Lu, and Yitong Yin. Correlation decay up to uniqueness in spin systems. In *Proc. of SODA'13*, pages 67–84, 2013. doi:10.1137/1.9781611973105.5.

40 Fabio Martinelli. Lectures on Glauber dynamics for discrete spin models. In *Lectures on probability theory and statistics*, pages 93–191. Springer, 1999. doi:10.1007/978-3-540-48115-7_2.

41 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.

42 Marcus Michelen and Will Perkins. Analyticity for classical gasses via recursion. *CoRR*, abs/2008.00972, 2020. URL: http://arxiv.org/abs/2008.00972.

43 Sarat Babu Moka, Sandeep Juneja, and Michel R. H. Mandjes. Analysis of perfect sampling methods for hard-sphere models. *SIGMETRICS Performance Evaluation Review*, 45(3):69–75, 2018. doi:10.1145/3199524.3199536.

44 Elchanan Mossel and Allan Sly. Gibbs rapidly samples colorings of $g(n, d/n)$. *Probability Theory and Related Fields*, 148(1-2):37–69, 2010. doi:10.1007/s00440-009-0222-x.

45 Elchanan Mossel, Dror Weitz, and Nicholas Wormald. On the hardness of sampling independent sets beyond the tree threshold. *Probability Theory and Related Fields*, 143(3-4):401–439, 2009. doi:10.1007/s00440-007-0131-9.

46 Viresh Patel and Guus Regts. Deterministic polynomial-time approximation algorithms for partition functions and graph polynomials. *SIAM Journal on Computing*, 46(6):1893–1919, 2017. doi:10.1137/16M1101003.

47 Mathew D Penrose. Self-avoiding walks and trees in spread-out lattices. *Journal of Statistical Physics*, 77(1-2):3–15, 1994. doi:10.1007/BF02186829.

48 Ricardo Restrepo, Jinwoo Shin, Prasad Tetali, Eric Vigoda, and Linji Yang. Improved mixing condition on the grid for counting and sampling independent sets. *Probability Theory and Related Fields*, 156(1-2):75–99, 2013. doi:10.1007/s00440-012-0421-8.

49 Alistair Sinclair, Piyush Srivastava, Daniel Štefankovič, and Yitong Yin. Spatial mixing and the connective constant: Optimal bounds. *Probability Theory and Related Fields*, 168(1-2):153–197, 2017. doi:10.1007/s00440-016-0708-2.

50 Allan Sly. Computational transition at the uniqueness threshold. In *Proc. of FOCS'10*, pages 287–296, 2010. doi:10.1109/FOCS.2010.34.

51 J. van den Berg and R. Brouwer. Random sampling for the monomer–dimer model on a lattice. *Journal of Mathematical Physics*, 41(3):1585–1597, 2000. doi:10.1063/1.533198.

52 Juan C Vera, Eric Vigoda, and Linji Yang. Improved bounds on the phase transition for the hard-core model in 2-dimensions. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, pages 699–713. Springer, 2013. doi:10.1007/978-3-642-40328-6_48.

53 Dror Weitz. Combinatorial criteria for uniqueness of Gibbs measures. *Random Structures & Algorithms*, 27(4):445–475, 2005.

54 Dror Weitz. Counting independent sets up to the tree threshold. In *Proc. of STOC'06*, pages 140–149, 2006. doi:10.1145/1132516.1132538.