

Algorithms for hard-constraint point processes via discretization

Tobias Friedrich¹[0000–0003–0076–6308], Andreas Göbel¹[0000–0002–5180–7205], Maximilian Katzmann², Martin S. Krejca³[0000–0002–1765–1219], and Marcus Pappik¹[0000–0002–2480–3073]

¹ Hasso Plattner Institute, University of Potsdam, Potsdam, Germany

{tobias.friedrich, andreas.goebel, marcus.pappik}@hpi.de

² Karlsruhe Institute of Technology, Karlsruhe, Germany

maximilian.katzmann@kit.edu

³ Sorbonne University, CNRS, LIP6, Paris, France

martin.krejca@lip6.fr

Abstract We study the algorithmic applications of a natural discretization for the hard-sphere model and the Widom–Rowlinson model in a region of d -dimensional Euclidean space $\mathbb{V} \subset \mathbb{R}^d$. These continuous models are frequently used in statistical physics to describe mixtures of one or multiple particle types subjected to hard-core interactions. For each type, particles are distributed according to a Poisson point process with a type-specific activity parameter, called *fugacity*. The *Gibbs distribution* over all possible system states is characterized by the mixture of these point processes conditioned that no two particles are closer than some type-dependent distance threshold. A key part in better understanding the Gibbs distribution is its normalizing constant, called *partition function*. Our main algorithmic result is the first deterministic approximation algorithm for the partition function of the hard-sphere model and the Widom–Rowlinson model in box-shaped regions of Euclidean space. Our algorithms have quasi-polynomial running time in the volume of the region $\nu(\mathbb{V})$ if the fugacity is below a certain threshold. For the d -dimensional hard-sphere model with particles of unit volume, this threshold is $e/2^d$. As the number of dimensions d increases, this bound asymptotically matches the best known results for randomized approximation of the hard-sphere partition function. We prove similar bounds for the Widom–Rowlinson model. To the best of our knowledge, this is the first rigorous algorithmic result for this model.

Keywords: partition function · hard-sphere model · Widom–Rowlinson model · deterministic and randomized approximation · sampling

1 Introduction

Statistical physics models complex systems of interacting particles as probability distributions. The main goal is to explain the macroscopic properties of such a *spin system* when it is only described by the microscopic interactions among its

particles. Two classical models in this area are the hard-sphere model, central in the analysis of thermodynamics of liquids and liquid mixtures [11,3], and the Widom–Rowlinson model, which explains the evaporation of liquids [25]. These models have in common that each of their states consists of a finite set of points in a bounded region $V \subset \mathbb{R}^d$ of d -dimensional Euclidean space, distributed according to a Poisson point process of some intensity $\lambda \in \mathbb{R}_{\geq 1}$. The parameter λ is usually called *fugacity*. The points correspond to centers of spherical particles of some radius $r \in \mathbb{R}_{\geq 0}$ that are constrained by hard-core interactions, i.e., two particles are not allowed to occupy the same space. The probability distribution characterizing the valid states of a model is called its *Gibbs distribution*. A key part in better understanding this distribution is its normalizing factor: the *partition function* of the model, which is formally defined as a weighted integral over all valid configurations of the system.

Computing this partition function and sampling from the Gibbs distribution are the two main algorithmic tasks that are related to such models. These computational problems have been studied extensively for several decades [16,15,13,10,8,17], resulting in a variety of new techniques, such as the Monte Carlo method [16]. Whereas initially the focus was on the sampling problem, recently rigorous results on computing the partition function appeared more often [8,17,7]. Since computing such partition functions is notoriously difficult [23], they are typically approximated. Specifically, the goal is to determine the fugacity regime such that there is a relative ε -approximation algorithm for the partition function that runs efficiently in the volume of the considered region $\nu(V)$ and ε^{-1} . State-of-the-art results in that regard are usually achieved by randomized algorithms [8,17,7]. For example for the hard-sphere model, the best known bound is $\lambda < e/((1 - 1/8^{d+1})2^d \nu(\mathbb{B}(r)))$, where r is the radius of a single particle and $\nu(\mathbb{B}(r))$ is the volume of a ball of radius r . The result was obtained by Michelen and Perkins [17] using a Markov chain Monte Carlo approach, which is inherently probabilistic. Significantly less is known in the setting of deterministic approximation. The question of efficient deterministic approximation in a discrete computational model is especially elusive due to the fact that the space of possible point configurations of the model is continuous. It is even unclear how to obtain an exponential time brute-force algorithm for approximating the partition function. This raises the question if an efficient deterministic approximation exists for a comparable parameter regime.

Our contributions We propose a deterministic approximation algorithm for the hard-sphere model and the Widom–Rowlinson model on box-shaped regions $V \subset \mathbb{R}^d$ of Euclidean space. Our algorithms have quasi-polynomial running time for a fugacity regime that is comparable with the best known bounds for randomized approximations. Specifically, our approach applies to the hard-sphere model for $\lambda < e/(2^d \nu(\mathbb{B}(r)))$. Note that, as the number of dimensions d increases, this becomes equivalent to the best known randomized result. For the Widom–Rowlinson model with $q \in \mathbb{N}_{\geq 1}$ particle types, each with radius bounded by some $r \in \mathbb{R}$, our algorithm applies if the fugacity of each particle type is bounded by

$\lambda < e/((q-1)2^d\nu(\mathbb{B}(r)))$. To the best of our knowledge, this is the first rigorous computational result for approximating the partition function of the continuous Widom–Rowlinson model.

We obtain our results by viewing the hard-sphere model and the Widom–Rowlinson model under a common framework, which we call hard-constraint point processes. We then study a natural method to discretize such hard-constraint point processes. More precisely, we give sufficient conditions such that the partition function of a hard-constraint point process is closely approximated by the partition function of a discrete hard-core model on a geometric graph, based on a finite point set $X \subset \mathbb{V}$. This generalizes our previous discretization methods for hard-sphere models in box-shaped regions [8]. If X is of size polynomial in the volume $\nu(\mathbb{V})$, our results follow immediately from known approximation algorithms for the discrete hard-core model. For box-shaped regions $\mathbb{V} \in [0, \ell]^d$, we give a suitable point set $X \subset \mathbb{V}$ of size $O(\nu(\mathbb{V})^2)$, which greatly improves previously known super-exponential bounds [8]. We further argue that this bound is tight.

If we allow for randomization, we can use the same discretization scheme to give a fully polynomial randomized approximation for the partition functions of the hard-sphere model and the Widom–Rowlinson model for the same parameter regime to which our deterministic results apply. For the hard-sphere model, this simplifies the algorithm given in [8] and, regarding the Widom–Rowlinson model, this is the first fully polynomial approximation result in this regime. Assuming a continuous model of computation, we further show that our discretization can be used to obtain an approximate sampler for both models in the same fugacity regime.

1.1 Hard-constraint point processes

As discussed earlier, the hard-sphere model and the Widom–Rowlinson model have in common that they are characterized by a point process with hard-core interactions in a region \mathbb{V} . To remove redundancy, we define a more general class of models that includes both of the above. We refer to this class as *hard-constraint point processes*. Once we establish our results for this class of models, the corresponding statements for the hard-sphere model and the Widom–Rowlinson model follow immediately.

Let $\mathbb{V} \subset \mathbb{R}^d$ be bounded and measurable, and let $q \in \mathbb{N}_{\geq 1}$. The model represents the distribution of particles of q types, labeled by elements in $[q] := [1, q] \cap \mathbb{N}$, on \mathbb{V} . Particles of the same types are assumed to be indistinguishable. Let $\lambda: [q] \rightarrow \mathbb{R}_{\geq 0}$ be a function that equips each particle type with a fugacity. For each particle type $i \in [q]$, we assume that the positions of particles are distributed according to a (labeled) Poisson point process of intensity $\lambda(i)$ on \mathbb{V} . To add the constraints to the model, let $R \in \mathbb{R}_{\geq 0}^{q \times q}$ be a symmetric $q \times q$ matrix, called the *interaction matrix*. We condition the mixture of point processes by rejecting all configurations that contain particles at positions $x_1, x_2 \in \mathbb{V}$ with corresponding particle types $\tau_1, \tau_2 \in [q]$ with distance $d(x_1, x_2) < R(\tau_1, \tau_2)$. That is, the entries of R determine the minimum distance that particles of the respective types can have. Especially note that $R(\tau_1, \tau_2) = 0$ means that the particle types $\tau_1, \tau_2 \in [q]$

are not subjected to any pairwise constraints. Thus, if all entries of R are set to 0, we obtain a simple mixture of Poisson point processes.

For any instance (\mathbb{V}, R, λ) of a hard-constraint point process, the above characterizes a *Gibbs distribution* as follows. For all $k \in \mathbb{N}$ and all type assignments $\tau: [k] \rightarrow [q]$, let $D_\tau^{(R)}: (\mathbb{R}^d)^k \rightarrow \{0, 1\}$ be the function that indicates for a tuple of particles positions $\mathbf{x} = (x_i)_{i \in [k]} \in (\mathbb{R}^d)^k$ whether it forms a valid configuration, assuming that for each $i \in [k]$, the particle at x_i is of type $\tau(i)$. Formally,

$$D_\tau^{(R)}(\mathbf{x}) = \prod_{i, j \in [k]: i < j} \mathbb{1}\{d(x_i, x_j) \geq R(\tau(i), \tau(j))\},$$

noting that $D_\tau^{(R)}(\mathbf{x}) = 1$ for the case that $k = 0$. For each tuple (\mathbf{x}, τ) as above, the Gibbs distribution of (\mathbb{V}, R, λ) is defined via the probability density

$$\mu^{(\mathbb{V}, R, \lambda)}(\mathbf{x}, \tau) = \frac{\frac{1}{k!} \left(\prod_{i \in [k]} \lambda(\tau(i)) \right) D_\tau^{(R)}(\mathbf{x})}{Z(\mathbb{V}, R, \lambda)},$$

where the normalizing constant $Z(\mathbb{V}, R, \lambda)$ is called the *partition function*:

$$Z(\mathbb{V}, R, \lambda) = 1 + \sum_{k \in \mathbb{N}_{\geq 1}} \frac{1}{k!} \sum_{\tau: [k] \rightarrow [q]} \left(\prod_{i \in [k]} \lambda(\tau(i)) \right) \int_{\mathbb{V}^k} D_\tau^{(R)}(\mathbf{x}) d\nu^{d \times k},$$

with $\nu^{d \times k}$ denoting the product of k Lebesgue measures, each on \mathbb{R}^d . We proceed by showing how the hard-sphere model and the Widom–Rowlinson model are recovered from this framework.

Hard-sphere model: The hard-sphere model describes the distribution of particles of a single type subjected to hard-core interactions. It is parameterized by a radius $r \in \mathbb{R}_{\geq 0}$ and a single fugacity $\lambda \in \mathbb{R}_{\geq 0}$. Informally speaking, it is a simple Poisson point process of intensity λ , conditioned on no two points being closer than $2r$ (i.e., when placing balls of radius r at each of the points, they must be non-overlapping). This is equivalent to a hard-constraint point process $(\mathbb{V}, R_{\text{HS}}, \lambda)$ with $q = 1$ particle types, where we abuse notation and treat λ as a constant function, and set R_{HS} to be a 1×1 matrix containing only the entry $2r$.

Widom–Rowlinson model: We consider the most general version of the Widom–Rowlinson model, although we might impose certain restrictions for some of our algorithmic results. The Widom–Rowlinson model describes the interaction of particles of $q \in \mathbb{N}_{\geq 1}$ types, each equipped with a radius $r_i \in \mathbb{R}_{\geq 0}$ and a fugacity $\lambda_i \in \mathbb{R}_{\geq 0}$ for $i \in [q]$. Informally speaking, the resulting distribution is a mixture of q Poisson point processes, each with its own intensity λ_i , with the condition that particles of the same type can be arbitrarily close to each other, but particles of different types $i, j \in [q]$ need to have a distance of at least $r_i + r_j$ (i.e., when placing a ball of radius r_i at each point of type $i \in [q]$, balls of different types must be non-overlapping). This is equivalent to a hard-constraint point process

$(\mathbb{V}, R_{\text{WR}}, \lambda)$ with q particle types, where we set $\lambda(i) = \lambda_i$ for all $i \in [q]$, and, for all $i, j \in [q]$,

$$R_{\text{WR}}(i, j) = \begin{cases} 0 & \text{if } i = j, \\ r_i + r_j & \text{otherwise.} \end{cases}$$

1.2 Reduction to a discrete hard-core model

We investigate a natural discretization to turn a (continuous) hard-constraint point process into a (discrete) hard-core model. For an undirected graph $G = (V, E)$ and a function $\lambda: V \rightarrow \mathbb{R}_{\geq 0}$, the *(multivariate) hard-core model* is defined by the tuple (G, λ) . Let $\mathcal{I}(G)$ denote the independent sets of G . We associate the hard-core model (G, λ) with a distribution $\mu_{\text{HC}}^{(G, \lambda)}$ on $\mathcal{I}(G)$ that assigns each independent set $I \in \mathcal{I}(G)$ a probability proportional to $\prod_{v \in I} \lambda(v)$, and we call $\mu_{\text{HC}}^{(G, \lambda)}$ the Gibbs distribution of (G, λ) . Its normalizing constant $Z_{\text{HC}}(G, \lambda)$ is called the partition function of (G, λ) . Formally, this is

$$Z_{\text{HC}}(G, \lambda) = \sum_{I \in \mathcal{I}(G)} \prod_{v \in I} \lambda(v) \text{ and, for } I \in \mathcal{I}(G), \text{ we have } \mu_{\text{HC}}^{(G, \lambda)}(I) = \frac{\prod_{v \in I} \lambda(v)}{Z_{\text{HC}}(G, \lambda)}.$$

Our main goal is to reduce the problem of approximating the partition function of a hard-constraint point process to approximating the partition function of a suitable hard-core model. The advantage of this approach is that computational properties of hard-core models were extensively studied by the computer science community. In the uniform case (i.e., λ is constant), celebrated results include the deterministic approximation of $Z_{\text{HC}}(G, \lambda)$ in time $|V|^{\text{O}(\ln(\Delta))}$ [24] and the randomized approximation of $Z_{\text{HC}}(G, \lambda)$ in time $\tilde{\text{O}}(|V|^2)$ [1, 22] for all graphs G of maximum degree Δ and all λ below the *tree threshold* $\lambda_c(\Delta) = \frac{(\Delta-1)^{\Delta-1}}{(\Delta-2)^\Delta}$. For general graphs, approximation above this threshold is NP-hard [21, 9].

We proceed by describing how we map a hard-constraint point process to a hard-core model. Given an instance of a hard-constraint point process (\mathbb{V}, R, λ) with $q \in \mathbb{N}_{\geq 1}$ particle types and a finite non-empty set of points $X \subset \mathbb{V}$, we construct a simple undirected graph $G_X = (V_X, E_X)$ as follows. For each point in X and each type, we have a vertex in V_X . Two distinct vertices are connected by an edge if and only if no two particles of the corresponding types are allowed to occupy the respective positions in a valid configuration. Consequently, independent sets of the graph correspond to valid particle configurations. We aim to “simulate” the original continuous model by a hard-core model on that graph.

Formally, we get the following construction:

- For each point $x \in X$ and each type $i \in [q]$, we construct a vertex $v_x^{(i)}$. Furthermore, for each $i \in [q]$, we set $V_X^{(i)} = \{v_x^{(i)} \mid x \in X\}$, and we define $V_X = \bigcup_{i \in [q]} V_X^{(i)}$.

- For each $i, j \in [q]$ and $x, y \in X$, we connect $v_x^{(i)}, v_y^{(j)} \in V_X$ with an edge in E_X if

$$\neg(x = y \text{ and } i = j) \text{ and } d(x, y) < R(i, j),$$

where the condition $\neg(x = y \text{ and } i = j)$ prevents self-loops.

Additionally, we define a function $\lambda_X: V_X \rightarrow \mathbb{R}_{\geq 0}$ such that, for all $i \in [q]$ and all $x \in X$, $\lambda_X(v_x^{(i)}) = \frac{\nu(\mathbb{V})}{|X|} \lambda(i)$. Note that for all $i \in [q]$ and all $x, y \in X$, this means that $\lambda_X(v_x^{(i)}) = \lambda_X(v_y^{(i)})$, and we sometimes abuse notation and write $\lambda_X(i)$ instead. We call (G_X, λ_X) the *hard-core representation* of (\mathbb{V}, R, λ) based on X . Our goal is to find conditions, such that the hard-core partition function $Z_{\text{HC}}(G_X, \lambda_X)$ closely approximates the partition function of the original point process $Z(\mathbb{V}, R, \lambda)$.

We restrict to what we call *canonical discretization* for box-shaped regions $\mathbb{V} = [0, \ell]^d$ with $\ell \in \mathbb{R}_{>0}$. Formally, the canonical point set is parameterized by a resolution $\rho \in \mathbb{R}_{>0}$. Let $\mathbb{G}_n = [0, n]^d \cap \mathbb{N}^d$ be a d -dimensional positive integer grid. For all box-shaped regions $\mathbb{V} = [0, \ell]^d$, $\ell \in \mathbb{R}_{>0}$, we call $\rho \in \mathbb{R}_{>0}$ a *feasible resolution* if and only if $\ell\rho \in \mathbb{N}$. Given a feasible resolution $\rho \in \mathbb{R}_{>0}$, we define the canonical point set for resolution ρ as

$$X_\rho = \frac{1}{\rho} \mathbb{G}_{\ell\rho} = \left\{ \left(x^{(1)}, \dots, x^{(d)} \right) \in \mathbb{V} \mid \forall i \in [d] \exists m \in \mathbb{N}: x^{(i)} = \frac{m}{\rho} \right\}.$$

Our main technical result is that the partition function of the hard-core representation based on a point set $X \subset \mathbb{V}$ closely approximates the partition function of the respective hard-constraint point process.

Theorem 1. *Let (\mathbb{V}, R, λ) be a hard-constraint point process with $q \in \mathbb{N}_{\geq 1}$ particle types and $\mathbb{V} = [0, \ell]^d$ for some $\ell \in \mathbb{R}_{>0}$. For all $\varepsilon_{\text{D}} \in (0, 1]$ there exists $\rho_{\varepsilon_{\text{D}}} \in \Theta\left(\varepsilon_{\text{D}}^{-1/d} \nu(\mathbb{V})^{1/d}\right)$ such that for all feasible resolutions $\rho \geq \rho_{\varepsilon_{\text{D}}}$ we have*

$$e^{-\varepsilon_{\text{D}}} Z(\mathbb{V}, R, \lambda) \leq Z_{\text{HC}}(G_{X_\rho}, \lambda_{X_\rho}) \leq e^{\varepsilon_{\text{D}}} Z(\mathbb{V}, R, \lambda).$$

Since [Theorem 1](#) is at the core of our algorithmic results, we briefly overview its proof. Given our point set X_ρ , we define an allocation function $\Phi: \mathbb{V} \rightarrow X_\rho$ that maps each point $y = (y^{(1)}, \dots, y^{(d)}) \in \mathbb{V}$ to the closest point $x = (x^{(1)}, \dots, x^{(d)}) \in X_\rho$ such that $x^{(i)} \leq y^{(i)}$ for all dimensions $i \in [d]$. The main idea is to compare the contribution of points in $\Phi^{-1}(x)$ to the partition function $Z(\mathbb{V}, R, \lambda)$ with the contribution of the vertex of G_{X_ρ} that corresponds to x to $Z_{\text{HC}}(G_{X_\rho}, \lambda_{X_\rho})$. This comparison needs to consider two types of errors between the two partition functions, which we bound. The first one is due to the Poisson point process of the continuous model potentially generating more than one point in $\Phi^{-1}(x)$, for some $x \in X_\rho$, which cannot be represented as part of the hard-core partition function. To bound this error, we reduce the continuous model to an intermediate multiset version of the hard-core model, where valid configurations are allowed to contain multiple copies of the same vertex. Then,

we bound the difference between the partition function of the multiset hard-core model and $Z_{\text{HC}}(G_{X_\rho}, \lambda_{X_\rho})$. The second error type is caused by the fact that Φ potentially allocates configurations that are valid for the hard-constraint point process to configurations that are invalid for the hard-core model and vice versa. This happens because, e.g., two points $x_1, x_2 \in \mathbb{V}$ of type τ_1, τ_2 , respectively, with $d(x_1, x_2) \geq R(\tau_1, \tau_2)$ get mapped to $\Phi(x_1), \Phi(x_2) \in X$ with $d(\Phi(x_1), \Phi(x_2)) < R(\tau_1, \tau_2)$. The choice of $\rho_{\varepsilon_{\text{D}}}$ ensures that each point $x \in \mathbb{V}$ is allocated to a point $\Phi(x) \in X_\rho$ with distance $d(x, \Phi(x)) \leq \varepsilon_{\text{D}}$. This allows us to bound this error by the difference of two partition functions with slightly different values in their exclusion matrices R . In Euclidean space, we can express this as a difference between partition functions with the original exclusion matrix R but on differently scaled versions of the regions $(1 \pm \alpha)\mathbb{V}$ for some small α . Note that $(1 - \alpha)\mathbb{V} \subseteq \mathbb{V} \subseteq (1 + \alpha)\mathbb{V}$ and using basic properties of $Z(\mathbb{V}, R, \lambda)$, such as monotonicity and log-subadditivity in \mathbb{V} , finally allows us to rewrite that error multiplicatively in $Z(\mathbb{V}, R, \lambda)$, which concludes the proof of [Theorem 1](#).

For the case of the hard-sphere model on box-shaped regions, the same discretization scheme was used before to obtain a randomized approximation algorithm [\[8\]](#). The main difference between the proof of [Theorem 1](#) and discretization result given in [\[8\]](#) is that we bound the discretization error by translating it to the difference of partition functions of hard-constraint point processes on scaled regions $(1 \pm \alpha)\mathbb{V}$. In contrast to that, the bound in [\[8\]](#) was derived by considering the error based on pairs of particles, leading to a bound on the required number of vertices that is super-exponential in $\nu(\mathbb{V})$. To deal with this exponential graph size, the previous algorithm needs to utilize the succinct representation and structural information of the produced graph, which was highly specific to the hard-sphere model. Since our refined bound yields much smaller graphs, we are able to use the results on the hard-core model as out-of-the-box algorithms. Besides the aforementioned algorithmic simplifications, our result extends the applicability to other hard-constraint point processes including the Widom–Rowlinson model. Furthermore, it allows for the first efficient deterministic approximation for the partition functions of both models in a fugacity regime that is comparable with the best known randomized results.

1.3 Approximation algorithms via canonical discretization

We show how [Theorem 1](#) is used to obtain approximations for the partition functions of the hard-sphere model and the Widom–Rowlinson model. To this end, we develop our result in the setting of general hard-constraint point processes and obtain the model-specific results as corollaries.

We first study the case where the considered hard-constraint point process only contains a single particle type (e.g., the hard-sphere model) or where all particle types have the same fugacity (e.g., the Widom–Rowlinson model with uniform fugacities). In this case, the hard-core representation is a uniform hard-core model, i.e., there is a constant λ such that, for each vertex v , it holds that $\lambda(v) = \lambda$. A selection of algorithmic results are applicable in this setting if the hard-core model in question is below the tree threshold. We obtain our algorithms by bounding the

maximum degree of G_X and using the known approximation algorithms for the hard-core model. Namely, using the deterministic approximation algorithm for the hard-core partition function introduced by Weitz [24], we obtain a deterministic approximation of the partition functions of hard-constraint point processes with running time quasi-polynomial in the volume of \mathbb{V} . Furthermore, we obtain fully polynomial randomized approximations in the same parameter regime using a Markov chain Monte Carlo algorithm, for which polynomial running time was recently proven in a sequence of papers [2,5,6,4,1].

In order to formally state our algorithmic results for the class of hard-constraint point processes, we introduce the *volume exclusion matrix* $B \in \mathbb{R}_{\geq 0}^{q \times q}$ of a hard-constraint point process (\mathbb{V}, R, λ) with $q \in \mathbb{N}_{\geq 1}$ particles types. For all $i, j \in [q]$, we set $B(i, j) = \nu(\mathbb{B}(R(i, j)))$, where $\mathbb{B}(r)$ denotes a d -dimensional ball of radius r . Intuitively, the entry $B(i, j)$ gives an upper bound on the volume of the region around a particle of type $i \in [q]$ in which no particle of type $j \in [q]$ can be placed.

Our main algorithmic result for general hard-constraint point processes is the following.

Theorem 2. *Let (\mathbb{V}, R, λ) be a hard-constraint point process with $q \in \mathbb{N}_{\geq 1}$ particle types, $\mathbb{V} = [0, \ell]^d$ for some $\ell \in \mathbb{R}_{>0}$, and let λ be a constant. Let B be the corresponding volume exclusion matrix and denote its L_1 -norm by $\|B\|_1$. If*

$$\lambda < \frac{e}{\|B\|_1},$$

then for all $\varepsilon_A \in (0, 1]$, there is a quasi-polynomial deterministic ε_A -approximation algorithm for $Z(\mathbb{V}, R, \lambda)$ with running time $\left(\frac{\nu(\mathbb{V})}{\varepsilon_A}\right)^{\Theta(\ln(\nu(\mathbb{V})/\varepsilon_A))}$ and a randomized ε_A -approximation algorithm for $Z(\mathbb{V}, R, \lambda)$ with running time $\tilde{O}(\nu(\mathbb{V})^4 \varepsilon_A^{-2})$.

We derive our algorithmic results for the hard-core model and the Widom–Rowlinson model with uniform fugacities and radii from [Theorem 2](#). For the hard-sphere model, we obtain the following corollary.

Corollary 1. *Let $\mathbb{V} = [0, \ell]^d$ for some $\ell \in \mathbb{R}_{>0}$. Further, let $r \in \mathbb{R}_{>0}$ and $\lambda \in \mathbb{R}_{\geq 0}$. Denote by $Z_{\text{HS}}(\mathbb{V}, r, \lambda)$ the hard-sphere partition function on \mathbb{V} with particles of radius r and fugacity λ . If*

$$\lambda < \frac{e}{2^d \nu(\mathbb{B}(r))},$$

then for all $\varepsilon_A \in (0, 1]$, there is a quasi-polynomial deterministic ε_A -approximation algorithm for $Z_{\text{HS}}(\mathbb{V}, r, \lambda)$ with running time $\left(\frac{\nu(\mathbb{V})}{\varepsilon_A}\right)^{\Theta(\ln(\nu(\mathbb{V})/\varepsilon_A))}$ and a randomized ε_A -approximation algorithm for $Z_{\text{HS}}(\mathbb{V}, r, \lambda)$ with running time $\tilde{O}(\nu(\mathbb{V})^4 \varepsilon_A^{-2})$.

The most relevant point of the above statement is the quasi-polynomial deterministic approximation algorithm. It is the first deterministic result for a parameter regime that is comparable with recent randomized approaches and partially answers an open question of Michelen and Perkins [17].

For the Widom–Rowlinson model, [Theorem 2](#) yields the following approximation result.

Corollary 2. *Let $\mathbb{V} = [0, \ell]^d$, $\ell \in \mathbb{R}_{>0}$, let and $r \in \mathbb{R}_{>0}$ and $\lambda \in \mathbb{R}_{\geq 0}$. Let $Z_{\text{WR}}(\mathbb{V}, r, \lambda)$ denote the Widom–Rowlinson partition function on \mathbb{V} with $q \in \mathbb{N}_{\geq 1}$ particle types, each of radius r and fugacity λ . If*

$$\lambda < \frac{e}{(q-1)2^d \nu(\mathbb{B}(r))},$$

then for all $\varepsilon_A \in (0, 1]$, there is a quasi-polynomial deterministic ε_A -approximation algorithm for $Z_{\text{WR}}(\mathbb{V}, r, \lambda)$ with running time $\left(\frac{\nu(\mathbb{V})}{\varepsilon_A}\right)^{\Theta(\ln(\nu(\mathbb{V})/\varepsilon_A))}$ and a randomized ε_A -approximation algorithm for $Z_{\text{WR}}(\mathbb{V}, r, \lambda)$ with running time $\tilde{O}(\nu(\mathbb{V})^4 \varepsilon_A^{-2})$.

To the best of our knowledge, no other efficient approximation algorithm is known for this parameter regime. In fact, [Corollary 2](#) also applies to the non-uniform Widom–Rowlinson model when replacing λ and r with the maximum fugacity and maximum radius among all particle types.

1.4 Sampling via random perturbations

So far, we only discussed approximation results for the partition functions of the hard-sphere model and the Widom–Rowlinson model, although, in discrete spin systems, approximation algorithms often go hand in hand with sampling algorithms [\[14\]](#). Unfortunately, for continuous spin systems, a natural barrier is that outputting a sample, i.e., a tuple of points that represents a valid configuration, requires infinite floating-point precision. Thus, assuming a discrete computational model, as is common in computer science, no (approximate) sampling algorithm with meaningful error bounds in terms of total-variation distance can be obtained.

Assuming a computational model that performs arithmetic operations of floating-point values with arbitrary precision and can uniformly sample a random floating-point number from an interval, we use our discretization G_X to recover an approximate sampler for $\mu^{(\mathbb{V}, R, \lambda)}$. In practice, the floating-point precision of common discrete computational models might be seen as sufficient for applying our sampling approach.

Again, we focus on box-shaped regions $\mathbb{V} = [0, \ell]^d$ with canonical discretization X_ρ for some feasible resolution $\rho \in \mathbb{R}_{>0}$. Recall the definition of the allocation function $\Phi : \mathbb{V} \rightarrow X_\rho$. Our sampling algorithm first samples an independent set I from the Gibbs distribution of the hard-core model on $(G_{X_\rho}, \lambda_{X_\rho})$, using a known Markov chain method (see, e.g., [\[2,5,4,1\]](#)). Recall that each vertex $v_x^{(i)}$ of G_{X_ρ} corresponds to a point $x \in X_\rho$ of a particle of type i . Then, for each vertex $v_x^{(i)}$ in I , our sampler chooses a position in $\Phi^{-1}(x)$ uniformly at random and places a particle of type i at this position. This way, we obtain the following sampling analogue of [Theorem 2](#).

Theorem 3. *Let (\mathbb{V}, R, λ) be a hard-constraint point process with $q \in \mathbb{N}_{\geq 1}$ particle types, $\mathbb{V} = [0, \ell]^d$ for some $\ell \in \mathbb{R}_{>0}$, and let λ be a constant. Let B be the corresponding volume exclusion matrix and denote its L_1 -norm by $\|B\|_1$. If*

$$\lambda < \frac{e}{\|B\|_1},$$

then for all $\varepsilon_S \in (0, 1]$, there is an ε_S -approximate sampler for $\mu^{(\mathbb{V}, R, \lambda)}$ with running time $\tilde{O}(\nu(\mathbb{V})^2 \varepsilon_S^{-1})$.

Our sampling procedure does not always result in a valid configuration for the hard-constraint point process. However, by bounding the total-variation distance between the two distributions, $\mu^{(\mathbb{V}, R, \lambda)}$ and the distribution of our sampler, we bound the probability of the sampler returning an invalid configuration by some small constant. In this case, we reject this configuration and repeat the sampling.

1.5 Discussion and future directions

Our algorithmic results are centered around discretizations based on the hard-core model. For the hard-sphere model, a slightly better fugacity bound has recently been obtained without discretization by directly applying a Markov chain Monte Carlo method to the continuous problem, assuming a continuous model of computation [17]. This raises several questions.

First of all, note that the result by [17] is purely probabilistic. Thus, it would be interesting to see if the same fugacity bound can also be achieved for deterministic approximation. An obvious idea is to use more detailed insights about the properties of the graphs that result from the discretization. The result of [17] is obtained by studying the potential-weighted connective constant, which is a generalization of the connective constant on graphs. In the discrete setting, the connective constant was already used to improve algorithmic results for the hard-core model on certain graph classes [20]. Therefore, a promising candidate could be to investigate the connective constant of the resulting graphs. However, arguments were made that for the canonical discretization of the hard-sphere model, the connective constant is asymptotically equivalent to the maximum degree of the graph (see [19] and [8, Section 1.4]). Thus, a more sophisticated structural property might be required.

Second, it would be interesting to see if a deterministic approximation can also be obtained by directly working with the continuous model, instead of discretizing it. One approach could be to approximate the logarithm of the partition function instead, for example by using the cluster expansion. This has successfully been done in the setting of discrete models [18, 12]. However, in contrast to the discrete setting, we are not aware of any method to compute the coefficients of the expansion efficiently, as this again involves evaluating non-trivial multi-dimensional integrals. The challenge would be to evaluate these integrals with sufficient accuracy using a discrete deterministic algorithm.

Another interesting algorithmic question is whether the quasi-polynomial running time of the deterministic approximation algorithm can be improved to a polynomial. The two techniques yielding deterministic algorithms for the hard-core model have a running time of $n^{O(\log \Delta)}$ for graphs of maximum degree Δ , which corresponds in our setting to the quasi-polynomial running time. In the correlation decay method of Weitz [24], this running time comes from computing the self-avoiding walk tree of G . In Barvinok's interpolation method, this comes from enumerating induced connected subgraphs [18]. Improving the running time

dependency on the maximum degree Δ for deterministic algorithms on general graphs, if possible, seems to require significant conceptual insight. However, one could hope that the symmetric structure of G_X , produced by our discretization, leads to faster running times for the computational tasks used in any of these two algorithmic techniques.

Finally, it is worth pointing out that there are several models in statistical physics that do not fit into the framework we study. In contrast, the Markov chain Monte Carlo approach in [17] yields a randomized approximating for partition functions of Gibbs point processes with finite-range repulsive potentials, including the hard-sphere model as a special case. This approach is probabilistic, raising the question of deterministic algorithms for such general models.

Acknowledgments

Andreas Göbel was funded by the project PAGES (project No. 467516565) of the German Research Foundation (DFG). This project has received funding from the European Union’s Horizon 2020 research and innovation program under the Marie Skłodowska-Curie grant agreement No. 945298-ParisRegionFP and was partially funded by the HPI Research School on Data Science and Engineering.

References

1. Anari, N., Jain, V., Koehler, F., Pham, H.T., Vuong, T.D.: Entropic independence II: Optimal sampling and concentration via restricted modified log-Sobolev inequalities. CoRR **abs/2111.03247** (2021), <https://arxiv.org/abs/2111.03247>
2. Anari, N., Liu, K., Gharan, S.O.: Spectral independence in high-dimensional expanders and applications to the hardcore model. CoRR **abs/2001.00303** (2020), <http://arxiv.org/abs/2001.00303>
3. Boublik, T., Nezbeda, I., Hlavaty, K.: Statistical thermodynamics of simple liquids and their mixtures. Fundamental Studies in Engineering, Elsevier (1980)
4. Chen, X., Feng, W., Yin, Y., Zhang, X.: Rapid mixing of Glauber dynamics via spectral independence for all degrees. CoRR **abs/2105.15005** (2021), <http://arxiv.org/abs/2105.15005>
5. Chen, Z., Liu, K., Vigoda, E.: Rapid mixing of Glauber dynamics up to uniqueness via contraction. CoRR **abs/2004.09083** (2020), <http://arxiv.org/abs/2004.09083>
6. Chen, Z., Liu, K., Vigoda, E.: Optimal mixing of Glauber dynamics: Entropy factorization via high-dimensional expansion. In: STOC. pp. 1537–1550. ACM (2021), <https://doi.org/10.1145/3406325.3451035>
7. Friedrich, T., Göbel, A., Katzmann, M., Krejca, M., Pappik, M.: Using random graphs to sample repulsive gibbs point processes with arbitrary-range potentials. CoRR **abs/2204.01793** (2022), <http://arxiv.org/abs/2204.01793>
8. Friedrich, T., Göbel, A., Krejca, M., Pappik, M.: A spectral independence view on hard spheres via block dynamics. In: Proceedings of the 48th International Colloquium on Automata, Languages, and Programming (ICALP). vol. 198, pp. 66:1–66:15 (2021). <https://doi.org/10.4230/LIPIcs.ICALP.2021.66>
9. Galanis, A., Ge, Q., Stefankovic, D., Vigoda, E., Yang, L.: Improved inapproximability results for counting independent sets in the hard-core model. Random Structures & Algorithms **45**(1), 78–110 (2014). <https://doi.org/10.1002/rsa.20479>

10. Guo, H., Jerrum, M.: Perfect simulation of the hard disks model by partial rejection sampling. CoRR **abs/1801.07342** (2018), <https://arxiv.org/abs/1801.07342v2>
11. Hansen, J.P., McDonald, I.R.: Theory of Simple Liquids. Academic Press, 4th edn. (2013). <https://doi.org/10.1016/B978-0-12-387032-2.00013-1>
12. Helmuth, T., Perkins, W., Regts, G.: Algorithmic Pirogov–Sinai theory. In: Proceedings of the 51st Annual ACM Symposium on the Theory of Computing (STOC). pp. 1009–1020 (2019). <https://doi.org/10.1145/3313276.3316305>
13. Jenssen, M., Joos, F., Perkins, W.: On the hard sphere model and sphere packings in high dimensions. Forum of Mathematics, Sigma **7**, e1 (2019). <https://doi.org/10.1017/fms.2018.25>
14. Jerrum, M., Valiant, L.G., Vazirani, V.V.: Random generation of combinatorial structures from a uniform distribution. Theoretical Computer Science **43**, 169–188 (1986). [https://doi.org/10.1016/0304-3975\(86\)90174-X](https://doi.org/10.1016/0304-3975(86)90174-X)
15. Kendall, W.S.: Perfect simulation for the area-interaction point process. In: Probability Towards 2000, pp. 218–234. Springer (1998). https://doi.org/10.1007/978-1-4612-2224-8_13
16. Metropolis, N., Rosenbluth, A.W., Rosenbluth, M.N., Teller, A.H., Teller, E.: Equation of state calculations by fast computing machines. The Journal of Chemical Physics **21**(6), 1087–1092 (1953). <https://doi.org/10.1063/1.1699114>
17. Michelen, M., Perkins, W.: Strong spatial mixing for repulsive point processes. CoRR **abs/2202.08753** (2022), <https://arxiv.org/abs/2202.08753>
18. Patel, V., Regts, G.: Deterministic polynomial-time approximation algorithms for partition functions and graph polynomials. SIAM Journal on Computing **46**(6), 1893–1919 (2017). <https://doi.org/10.1137/16M1101003>
19. Penrose, M.D.: Self-avoiding walks and trees in spread-out lattices. Journal of Statistical Physics **77**(1-2), 3–15 (1994). <https://doi.org/10.1007/BF02186829>
20. Sinclair, A., Srivastava, P., Štefankovič, D., Yin, Y.: Spatial mixing and the connective constant: Optimal bounds. Probability Theory and Related Fields **168**, 153–197 (2017). <https://doi.org/10.1007/s00440-016-0708-2>
21. Sly, A.: Computational transition at the uniqueness threshold. In: Proceedings of the 51st Annual IEEE Symposium on Foundations of Computer Science (FOCS). pp. 287–296 (2010). <https://doi.org/10.1109/FOCS.2010.34>
22. Štefankovič, D., Vempala, S., Vigoda, E.: Adaptive simulated annealing: A near-optimal connection between sampling and counting. Journal of the ACM **56**(3), 1–36 (2009). <https://doi.org/10.1145/1516512.1516520>
23. Stoyan, D., Kendall, W.S., Chiu, S.N., Mecke, J.: Stochastic geometry and its applications. John Wiley & Sons (2013). <https://doi.org/10.1002/9781118658222>
24. Weitz, D.: Counting independent sets up to the tree threshold. In: Proceedings of the 38th Annual ACM Symposium on the Theory of Computing (STOC). pp. 140–149 (2006). <https://doi.org/10.1145/1132516.1132538>
25. Widom, B., Rowlinson, J.S.: New model for the study of liquid-vapor phase transitions. Journal of Chemical Physics **52**(4), 1670–1684 (1970). <https://doi.org/10.1063/1.1673203>