

Perfect Sampling for Hard Spheres from Strong Spatial Mixing

Konrad Anand*, Andreas Göbel†, Marcus Pappik‡, Will Perkins‡

Abstract

We provide a perfect sampling algorithm for the hard-sphere model on subsets of \mathbb{R}^d with expected running time linear in the volume under the assumption of strong spatial mixing. A large number of perfect and approximate sampling algorithms have been devised to sample from the hard-sphere model, and our perfect sampling algorithm is efficient for a range of parameters for which only efficient approximate samplers were previously known and is faster than these known approximate approaches. Our methods also extend to the more general setting of Gibbs point processes interacting via finite-range, repulsive potentials.

1 Introduction

Gibbs point processes, or classical gases, are mathematical models of interacting particles. In statistical physics they are used to model gases, fluids, and crystals, while in other fields they are used to model spatial phenomena such as the growth of trees in a forest, the distribution of stars in the universe, or the location of cities on a map (see e.g. [71, 62, 76, 12]).

Perhaps the longest and most intensively studied Gibbs point process is the hard-sphere model: a model of a gas in which the only interaction between particles is a hard-core exclusion in a given radius around each particle. That is, it is a model of a random packing of equal-sized spheres. Despite the simplicity of its definition, the hard-sphere model is expected to exhibit the qualitative behavior of a real gas [2], and in particular exhibit gas, liquid, and solid phases, thus giving evidence for the hypothesis, dating back to at least Boltzmann, that the macroscopic properties of a gas or fluid are determined by its microscopic interactions. This rich behavior exhibited by the hard-sphere model is very difficult to analyze rigorously, and the most fundamental questions about phase transition in this model are open mathematical problems [71, 52].

*Queen Mary, University of London

†Hasso Plattner Institute, University of Potsdam

‡School of Computer Science, Georgia Institute of Technology

In studying the hard-sphere model (or Gibbs point processes more generally), a fundamental task is to sample from the model. Sampling is used to estimate statistics, observe evidence of phase transitions, and perform statistical tests on data. A wide variety of methods have been proposed to sample from these distributions; for instance, the Markov chain Monte Carlo (MCMC) method was first proposed by Metropolis, Rosenbluth, Rosenbluth, Teller, and Teller [55] to sample from the two-dimensional hard-sphere model. Understanding sampling methods for point processes in theory and in practice is a major area of study [61, 62, 16, 39, 49], and advances in sampling techniques have led to advances in the understanding of the physics of these models [55, 2, 52, 6, 5, 16].

In this paper we will be concerned with provably efficient sampling from the hard-sphere model. Rigorous guarantees for sampling algorithms come in several different varieties. One question is what notion of ‘efficient’ to use; another is what guarantee we insist on for the output. In this paper we will provide an efficient sampling algorithm under the strictest possible terms with respect to both running time and accuracy of the output: a linear-time, *perfect* sampling algorithm.

For simplicity we focus on sampling from the hard-sphere model defined on finite boxes in \mathbb{R}^d . For fixed parameter values of the model, the typical number of points appearing in such a region is linear in the volume, and so any sampling algorithm will require at least this much time.

As for guarantees on the output, there are two main types of guarantees. The first type is an *approximate sampler*: the output of such an algorithm must be distributed within ε total variation distance of the desired target distribution. Perhaps the main approach to efficient sampling from distributions normalized by intractable normalizing constants is the MCMC method. In this approach, one devises a Markov chain with the target distribution as the stationary distribution and runs a given number steps of the chain from a chosen starting configuration; if the number of steps is at least the ε -mixing time, then the final state has distribution within ε total variation distance of the target [42, 68, 13]. In general, however, computing or bounding the mixing time can be a very challenging problem.

The second type of guarantee is that of a *perfect sampler* [66]. Such an algorithm has a running time that is random, but the distribution of the output is guaranteed to be *exactly* that of the target distribution. The main advantage of perfect sampling algorithms – and the primary reason they are studied and used in practice – is that one need not prove a theorem or understand the mixing time of a Markov chain to run the algorithm and get an accurate sample; one can simply run the algorithm and know that the output has the correct distribution. The drawback is that the running time may be very large, depending on the specific algorithm and on the parameter regime. Some naive sampling methods such as rejection sampling return perfect samples but are inefficient on large instances (exponential expected running time in the volume). The breakthrough of Propp and Wilson in introducing ‘coupling from the past’ [66, 67] was to devise a procedure for using a Markov chain transition matrix to design perfect sampling algorithms which, under some conditions, could run in time polylogarithmic in the size of a discrete state space (polynomial-time in the size of the graph of a spin system), matching the efficiency of fast

mixing Markov chains which only return approximate samples (see also [4, 51] for precedents in perfect sampling). The work of Propp and Wilson led to numerous constructions of perfect sampling algorithms for problems with both discrete and continuous state spaces including [17, 27, 45, 63, 28, 21, 46, 61, 23]. Notably, many of the first applications of Propp and Wilson’s technique were in designing perfect sampling algorithms for Gibbs point processes (though often without rigorous guarantees on the efficiency of the algorithms).

Perfect sampling continues to be a very active area of research today, with a special focus on improving the range of parameters for which perfect sampling algorithms can (provably) run in expected linear or polynomial time [8, 40, 30]

In this paper we design a perfect sampling algorithm for the hard-sphere model (and Gibbs point processes interacting with a finite-range, repulsive pair potential more generally) that is guaranteed to run in linear expected time for activity parameters up to the best known bound for efficient approximate sampling via MCMC.

What is this bound and how do we design the algorithm? One central theme in the analysis of discrete spin systems is the relationship between spatial mixing (correlation decay properties) and temporal mixing (mixing times of Markov chains) [35, 1, 75, 54, 15]. At a high level, these works show that for discrete lattice systems a strong correlation decay property (*strong spatial mixing*) implies a near-optimal convergence rate for local-update Markov chains like the Glauber dynamics. Recently it has been showed that strong spatial mixing in a discrete lattice model also implies the existence of efficient *perfect* sampling algorithms [18, 3]. In parallel, there has been work establishing the connection between strong spatial mixing and optimal temporal mixing for Markov chains in the setting of the hard-sphere model and Gibbs point processes [33, 57, 58]. At a high level, our aim is to combine these threads to show that strong spatial mixing for Gibbs point processes implies the existence of an efficient perfect sampler. One challenge is that the approaches of [18, 3] are inherently discrete in that key steps of the algorithms involve enumerating over all possible configurations in a subregion, something that is not possible in the continuum. To overcome this we make essential use of Bernoulli factories – a method for perfect simulation of a coin flip with a bias $f(p)$ given access to coin flips of bias p . Bernoulli factories have recently been used in perfect sampling algorithms for solutions to constraint satisfaction problems in [31, 32].

1.1 The hard-sphere model, strong spatial mixing, and perfect sampling

The hard-sphere model is defined on a bounded, measurable subset Λ of \mathbb{R}^d with an activity parameter $\lambda \geq 0$ that governs the density of the model and a parameter $r > 0$ that governs the range of interaction (though by re-scaling there is really only one meaningful parameter, and we could take $r = 1$ without loss of generality). In words, the hard-sphere model is the distribution of finite point sets in Λ obtained by taking a Poisson point process of activity λ on Λ and conditioning on the event that all pairs of points are at distance at least r from each other; in other words, on the event that spheres of radius $r/2$ centered at the given points form a sphere packing.

We can equivalently define the model more explicitly, and in doing so, introduce objects and notation we work with throughout the paper. To begin, let \mathcal{N}_Λ be the set of all finite point sets in Λ and write $\mathfrak{R}_\Lambda \subseteq 2^{\mathcal{N}}$ for the σ -field generated by the maps $\{\mathcal{N} \rightarrow \mathbb{N}_0, \eta \mapsto |\eta \cap B| \mid B \subseteq \Lambda \text{ Borel-measurable}\}$. The hard-sphere model (or in fact any Gibbs point process) is a probability measure μ_λ on the space $(\mathcal{N}_\Lambda, \mathfrak{R}_\Lambda)$.

Define for every $x_1, \dots, x_k \in \mathbb{R}^k$ the indicator that the points are centers of non-overlapping spheres of radius $r/2$; that is,

$$D(x_1, \dots, x_k) = \prod_{\{i,j\} \in \binom{[k]}{2}} \mathbb{1}_{\text{dist}(x_i, x_j) \geq r}.$$

Then define the partition function

$$Z_\Lambda(\lambda) = \sum_{k \geq 0} \frac{\lambda^k}{k!} \int_{\Lambda^k} D(x_1, \dots, x_k) dx_1 \dots dx_k.$$

For an event $A \in \mathfrak{R}_\Lambda$, the hard-sphere model assigns the probability

$$\mu_\lambda(A) = \frac{1}{Z_\Lambda(\lambda)} \sum_{k \geq 0} \frac{\lambda^k}{k!} \int_{\Lambda^k} \mathbb{1}_{\{x_1, \dots, x_k\} \in A} D(x_1, \dots, x_k) dx_1 \dots dx_k. \quad (1)$$

A very useful generalization of this model is to allow for a non-constant (but measurable) activity function $\lambda : \Lambda \rightarrow [0, \infty)$. Here the model is a Poisson process with inhomogenous activity λ conditioned on the points forming the centers of a sphere packing; the partition function is now

$$Z_\Lambda(\lambda) = \sum_{k \geq 0} \frac{1}{k!} \int_{\Lambda^k} \prod_{i=1}^k \lambda(x_i) D(x_1, \dots, x_k) dx_1 \dots dx_k$$

and the measure μ_λ is defined analogously to (1). This generalization allows modeling of non-homogenous spaces and generalizes the concept of imposing boundary conditions on the model. To see this, suppose we fix a particle configuration $\eta \in \mathcal{N}_\Lambda$ as boundary conditions. Additional points are forbidden within the balls of radius r around each point $x \in \eta$; we can implement the distribution of additional points by considering the measure μ_λ with $\lambda(y) = 0$ if $\text{dist}(y, x) < r$ for some $x \in \eta$; and $\lambda(y) = \lambda$ otherwise. We denote the resulting activity function by λ by λ_η . Further, we can use this generalization to restrict a point process to only place points in a subregion $\Lambda' \subseteq \Lambda$ by considering the measure $\mu_{\lambda \mathbb{1}_{\Lambda'}}$ with activity function $\lambda \mathbb{1}_{\Lambda'} : x \mapsto \lambda \mathbb{1}_{x \in \Lambda'}$. Of course the generalization to measurable activity functions is much more general than this, and activity functions λ need not be realizable by boundary conditions or restriction to a subregion, nor take only two values.

This generalization to activity functions is crucial for defining *strong spatial mixing*, the condition under which we can guarantee the efficiency of our perfect sampling algorithm.

To define the concept of strong spatial mixing we consider projections of the measure μ_λ to subregions $\Lambda' \subseteq \Lambda$. We write $\mu_\lambda[\Lambda']$ for the probability measure on $(\mathcal{N}_{\Lambda'}, \mathfrak{R}_{\Lambda'})$ induced by μ_λ (we make this definition formal in Section 3). We can impose two distinct boundary conditions on Λ' by choosing two different activity functions λ, λ' . Strong spatial mixing asserts that the distributions $\mu_\lambda[\Lambda'], \mu_{\lambda'}[\Lambda']$ are close in total variation when λ, λ' differ only on points far from Λ' ; i.e., when $\text{dist}(\Lambda', \text{supp}(\lambda - \lambda'))$ is large (as $\text{supp}(\lambda - \lambda')$ is the set of points at which the two activity functions disagree).

Writing $|\Lambda'|$ for the volume of Λ' , strong spatial mixing with exponential decay is defined as follows.

Definition 1.1. *Given $a, b \in \mathbb{R}_{>0}$, the hard-sphere model on \mathbb{R}^d exhibits (a, b) -strong spatial mixing up to $\lambda \in \mathbb{R}_{>0}$ if for all bounded measurable $\Lambda \subset \mathbb{R}^d$ the following holds: For all measurable $\Lambda' \subseteq \Lambda$ and all activity functions $\lambda, \lambda' \leq \lambda$ it holds that*

$$d_{TV}(\mu_\lambda[\Lambda'], \mu_{\lambda'}[\Lambda']) \leq a |\Lambda'| e^{-b \cdot \text{dist}(\Lambda', \text{supp}(\lambda - \lambda'))},$$

where $d_{TV}(\cdot, \cdot)$ denotes total variation distance.

This definition of strong spatial mixing comes from [58], which in turn adapted similar notions from discrete spin systems [15, 77]. Strong spatial mixing has proved to be an essential definition in the analysis, both probabilistic and algorithmic, of spin systems on graphs, and many recent works are focused on either proving strong spatial mixing for a particular model, range of parameters, and class of graphs (e.g. [77, 22, 53, 72, 69, 9]) or deriving consequences of strong spatial mixing (e.g. [73, 19, 50, 18, 3]).

Our main result is a linear expected-time perfect sampling algorithm for the hard-sphere model under the assumption of strong spatial mixing.

Theorem 1.2. *There is a perfect sampling algorithm for the hard-sphere model on finite boxes $\Lambda \subset \mathbb{R}^d$ with the property that if the hard-sphere model exhibits (a, b) -strong spatial mixing up to λ , then the expected running time of the algorithm at activity λ is $O(|\Lambda|)$, where the implied constant is a function of a, b , and λ .*

In particular, one can run the algorithm for any value of λ (without knowing whether or not strong spatial mixing holds) and the algorithm will terminate in finite time with an output distributed exactly as μ_λ ; under the assumption of strong spatial mixing the expected running time is guaranteed to be linear in the volume.

Using bounds from [58] on strong spatial mixing in the hard-sphere model, we obtain the following explicit bounds on the activities for which the algorithm is efficient.

Corollary 1.3. *The above perfect sampling algorithms runs in expected time $O(|\Lambda|)$ when $\lambda < \frac{e}{v_d(r)}$, where $v_d(r)$ is the volume of the ball of radius r in \mathbb{R}^d .*

In comparison, near-linear time MCMC-based approximate samplers were given in [58] for the same range of parameters (following results for more restricted ranges in [43, 33]). For perfect sampling from the hard-sphere model, linear expected time algorithms were given in [36, 25] for more restrictive ranges of parameters.

1.2 Gibbs point processes with finite-range repulsive potentials

We now give a closely related result in the more general setting of Gibbs point processes interacting via finite-range, repulsive pair potentials.

Gibbs point processes are defined via a density against an underlying Poisson point process. In general, this density is the exponential of (the negative of) an energy function on point sets that captures the interactions between points. In many of the most studied cases, this energy function takes a special form: it is the sum of potentials over pairs of points in a configuration.

A *pair potential* is a measurable symmetric function $\phi : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R} \cup \{\infty\}$. For a bounded, measurable activity function λ on Λ the Gibbs point process with pair potential ϕ on Λ is defined via the partition function

$$Z_\Lambda(\lambda) = \sum_{k \geq 0} \frac{1}{k!} \int_{\Lambda^k} \left(\prod_{i \in [k]} \lambda(x_i) \right) e^{-H(x_1, \dots, x_k)} dx_1 \dots dx_k$$

where

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

Again the corresponding probability measure μ_λ is obtained as in (1). A pair potential ϕ is *repulsive* if $\phi(x, y) \geq 0$ for all x, y . It is of *finite-range* if there exists $r \geq 0$ so that $\phi(x, y) = 0$ whenever $\text{dist}(x, y) > r$. As with the hard-sphere model, we can use the activity function to encode the influence of boundary conditions by defining the activity function $\lambda_\eta : y \mapsto \lambda e^{-\sum_{x \in \eta} \phi(x, y)}$ for any activity $\lambda \in [0, \infty)$ and particle configuration $\eta \in \mathcal{N}_\Lambda$. Moreover, strong spatial mixing for a Gibbs point process is defined exactly as in Definition 1.1.

The hard-sphere model is one example of a model interacting via a finite-range, repulsive pair potential; it is obtained by letting $\phi(x, y)$ take the value $+\infty$ if $\text{dist}(x, y) \leq r$ and 0 otherwise. The Strauss process [74, 44] is another such example.

Our next result is a near-linear expected time perfect sampling algorithm for Gibbs point processes interacting via finite-range, repulsive potentials under the assumption of strong spatial mixing.

Theorem 1.4. *Suppose ϕ is a finite-range, repulsive potential on \mathbb{R}^d and suppose ϕ exhibits (a, b) -strong spatial mixing up to λ for some constants $a, b > 0$. Then there is a perfect sampling algorithm for the Gibbs point process defined by ϕ and activity functions bounded by λ on boxes Λ in \mathbb{R}^d with expected running time $O(|\Lambda| \log^{O(1)} |\Lambda|)$.*

One difference between this algorithm and the hard-sphere algorithm of Theorem 1.2 is that this algorithm needs knowledge of the constants a, b in the assumption of strong spatial mixing, whereas the hard-sphere algorithm does not.

Using the results of [58], we can get explicit bounds for the existence of efficient perfect sampling algorithms in terms of the *temperedness constant* of the potential defined by

$$C_\phi := \sup_{x \in \mathbb{R}^d} \int_{\mathbb{R}^d} |1 - e^{-\phi(x,y)}| dy. \quad (2)$$

Under the assumption that ϕ is repulsive and of finite range r , we have $0 \leq C_\phi \leq v_d(r)$.

Corollary 1.5. *The above perfect sampling algorithm runs in expected time $O(|\Lambda| \log^{O(1)} |\Lambda|)$ when $\lambda < \frac{e}{C_\phi}$.*

Remark 1.6. *In fact, using the results of Michelen and Perkins [56], one can push the bound for strong spatial mixing up to e/Δ_ϕ , where $\Delta_\phi \leq C_\phi$ is the potential-weighted connective constant defined therein; our perfect sampling algorithm is efficient up to that point.*

1.3 Related work and future directions

Related work

In recent years there has been a moderate flurry of activity around proving rigorous results for Gibbs point processes in both the setting of statistical physics and probability theory and in the setting of provably efficient sampling algorithms.

Work on provably efficient approximate sampling methods for the hard-sphere model begins with the seminal paper of Kannan, Mahoney, and Montenegro [43], who used techniques from the analysis of discrete spin systems to prove mixing time bounds for Markov chains for the hard-sphere model. Improvements to the range of parameters for which fast mixing holds came in [29, 33], before Michelen and Perkins proved the bound $e/v_d(r)$ in [58], which we match with a perfect sampling algorithm in Corollary 1.3.

Perfect sampling algorithms for the hard sphere model have been considered in [27, 46, 21, 25, 38]. In terms of rigorous guarantees of efficiency, Huber proved a bound of $2/v_d(r)$ for a near-linear expected time perfect sampler in [36]. The perfect sampling algorithm of Guo and Jerrum in [25] does not match this bound, but the algorithm, based on ‘partial rejection sampling’ [26] is novel and particularly simple. Several of these approaches also apply for finite-range, repulsive potentials or can be extended to that setting (e.g. [60]).

In parallel, there has been much work on proving bounds on the range of activities for which no phase transition can occur in the hard-sphere model; and, in recent years in particular, the techniques used have close connections to algorithms and the study of Markov chains. The classic approach to proving absence of phase transition is by proving convergence of the cluster expansion; the original bound here is $1/(ev_d(r))$ due to Groeneveld [24]. In small dimensions (most significantly in dimension 2) improvements to the radius of convergence can be obtained [20]. On the other hand, this approach is inherently limited by the presence of non-physical singularities on the negative real axis. Alternative approaches avoiding this obstruction include using the equivalence of spatial and temporal mixing [33, 58]; or

disagreement percolation [10, 34, 7]. The best current bound for absence of phase transition for the hard-sphere model and for repulsive pair potentials is the bound of e/C_ϕ (and e/Δ_ϕ) obtained by Michelen and Perkins [57, 58, 56]. Theorem 1.4 brings the bound for efficient perfect sampling up to this bound.

On a technical level, the most relevant past work is [18], in which the authors prove that for discrete spins systems, strong spatial mixing and subexponential volume growth of a sequence of graphs imply the existence of an efficient perfect sampling algorithm. We take their approach as a starting point but need new ideas to replace their exhaustive enumeration of configurations.

A key step in our algorithm is the use of a Bernoulli factory to implement a Bayes filter. Bernoulli factories are algorithms by which a Bernoulli random variable with success probability $f(p)$ can be simulated (perfectly) by an algorithm with access to independent Bernoulli p random variables, where the algorithm does not know the value p . Whether a Bernoulli factory exists (and how efficient it can be) depends on the function $f(\cdot)$ and a priori bounds on the possible values p . Bernoulli factories have been studied in [64, 37, 14] and recently used in the design of perfect sampling algorithms for CSP solutions in [31, 32].

Future directions

There are a number of extensions and improvements to these results one could pursue. Perhaps most straightforward would be to relax the notion of strong spatial mixing from exponential decay to decay faster than the volume growth of \mathbb{R}^d and to extend the results to repulsive potentials of unbounded range but finite temperedness constant C_ϕ . Moreover, it would be nice to upgrade the guarantees of the algorithm in Theorem 1.4 to that of Theorem 1.2: that the algorithm does not need prior knowledge of the strong spatial mixing constants a, b to run correctly.

An ambitious and exciting direction would be to remove the assumption of a repulsive potential and find efficient perfect sampling algorithms for the class of *stable* potentials (see e.g. [65, 70, 71] for a definition). A stable potential is repulsive at short ranges but can include a weak attractive part; such potentials include the physically realistic Lenard-Jones potential among others [78]. This would require some very new ideas, as much of the recent probabilistic and algorithmic work on Gibbs point processes (e.g. [57, 58, 7, 41, 56]) has used repulsiveness as an essential ingredient (for one, repulsiveness of the potential implies stochastic domination by the underlying Poisson point process).

1.4 Outline of the paper

In Section 2, we describe the high-level idea and intuition behind the algorithm. In Section 3 we introduce some notation and present some preliminary results that we will use throughout the paper. In Section 4 we present the algorithm that we will apply to both hard spheres and more general processes. In Section 5 we prove correctness of the algorithm. In Section 6 we prove a technical lemma that will be crucial for showing efficiency of our algorithm under

the assumption of strong spatial mixing. In Section 7 we specialize to the hard-sphere model to complete the proof of Theorem 1.2. In Section 8 we work with finite-range, repulsive potentials to complete the proof of Theorem 1.4. In Section 9 we prove the running time bound for the Bernoulli factory used by our algorithm. The appendix contains some technical lemmas on measure theory and stochastic processes.

2 Intuitive idea behind the algorithm

Our algorithm is an adaptation to continuum models of the work by Feng, Guo, and Yin [18] on perfect sampling from discrete spin systems. We mimic their setting of a spin system on a graph $G = (V, E)$ by putting a graphical structure on sub-regions of our continuous space.

Let $\Lambda = [0, L]^d \subset \mathbb{R}^d$ be the region considered, $\lambda > 0$ the activity, and let ϕ be a repulsive potential of range $r > 0$. We subdivide Λ into $(\Lambda_{\mathbf{v}})_{\mathbf{v} \in \mathcal{V}}$, a set of smaller boxes of side length r indexed by vertices of a graph: each box corresponds to a vertex and boxes are connected if they are within r of each other, i.e., particles in the boxes can interact directly through the potential ϕ . We fix the index set for the boxes to be $\mathcal{V} \subset \mathbb{N}_0^d$, where each $\mathbf{v} \in \mathcal{V}$ corresponds to the box $\Lambda_{\mathbf{v}} = [v_1 r, (v_1 + 1)r) \times \cdots \times [v_d r, (v_d + 1)r)$. We extend this notation to sets of indices $S \subseteq \mathcal{V}$ by setting $\Lambda_S = \bigcup_{\mathbf{v} \in S} \Lambda_{\mathbf{v}}$. Further, we denote by $\mathbb{B}_k(\mathbf{v})$ the set of indices $\mathbf{w} \in \mathcal{V}$ with $\|\mathbf{v} - \mathbf{w}\|_{\infty} \leq k$. To shorten notation, we write $\partial S = (\bigcup_{\mathbf{v} \in S} \mathbb{B}_1(\mathbf{v})) \setminus S$ for the outer boundary of a set of boxes indexed by $S \subseteq \mathcal{V}$.

Our algorithm runs iteratively, keeping track of two random variables: a point configuration $X_t \in \mathcal{N}_{\Lambda}$ with $X_0 = \emptyset$, and a set of ‘incorrect’ boxes $\mathcal{U}_t \subseteq \mathcal{V}$ with $\mathcal{U}_0 = \mathcal{V}$. With each iteration t we maintain the following *invariant*: the partial configuration $X_t \cap (\Lambda_{\mathcal{U}_t})^c$ is distributed according to the projection of μ_{λ} to $(\Lambda_{\mathcal{U}_t})^c$ under the boundary condition $X_t \cap \Lambda_{\mathcal{U}_t}$. It follows that X_t is distributed according to μ_{λ} once we reach the state $\mathcal{U}_t = \emptyset$.

We proceed by sketching an iteration of the algorithm. An example for the involved subregions is given in Figure 1. Each iteration runs as follows:

1. We choose $\mathbf{u}_t \in \mathcal{U}_t$ uniformly at random and attempt to ‘repair’ it by updating X_t on a neighborhood of boxes $B = \{\mathbf{u}_t\} \cup (\mathbb{B}_{\ell}(\mathbf{u}_t) \setminus \mathcal{U}_t)$ for some *update radius* $\ell \in \mathbb{N}$.
2. We sample a Bayes filter F_t (i.e., a Bernoulli random variable) with probability depending on the potential ϕ , the activity λ , and the current point configuration X_t on $\Lambda_{\mathbf{u}_t}$ and $\Lambda_{\partial B}$.
3.
 - a) If $F_t = 1$, we set $\mathcal{U}_{t+1} = \mathcal{U}_t \setminus \{\mathbf{u}_t\}$ and we get X_{t+1} by updating X_t on Λ_B according to a projection of μ_{λ} conditioned on the boundary configuration $X_t \cap (\Lambda_B)^c$.
 - b) If $F_t = 0$, the configuration is unchanged and we add the boundary boxes to our ‘incorrect’ list, i.e., $X_{t+1} = X_t$ and $\mathcal{U}_{t+1} = \mathcal{U}_t \cup \partial B$.

We use the Bayes filter, as in [18], to remove bias from the resulting distribution. To give some intuition for its role, suppose we run a naive version of the algorithm where we




Figure 1: The box-shaped region $\Lambda \subset \mathbb{R}^2$ is divided into boxes of side length r (dotted lines). The boxes \mathcal{U}_t are bordered by bold black lines. For \mathbf{u}_t as given and update radius $\ell = 2$, the corresponding set B of boxes to be updated is indicated by the red hatched area (falling left to right). Its boundary boxes ∂B are shown as blue hatched area (rising left to right). The boxes in $H = (\mathcal{U}_t \cup B)^c$ are shown with gray background.

always update X_t on Λ_B as in step 3.a) above. Assuming the desired invariant holds after t iterations, this naive algorithm gives a bias to the distribution of X_{t+1} proportional to $\frac{Z_{\Lambda_B \setminus \{\mathbf{u}_t\}}(\lambda_{X_t \cap \Lambda_{\partial B \cup \{\mathbf{u}_t\}}})}{Z_{\Lambda_B}(\lambda_{X_t \cap \Lambda_{\partial B}})}$. We choose the Bayes filter such that, conditioned on $F_t = 1$, the bias term gets canceled. This suggests the choice

$$\mathbb{P}[F_t = 1 \mid X_t, \mathcal{U}_t, \mathbf{u}_t] = C(\mathcal{U}_t, \mathbf{u}_t, X_t) \cdot \frac{Z_{\Lambda_B}(\lambda_{X_t \cap \Lambda_{\partial B}})}{Z_{\Lambda_B \setminus \{\mathbf{u}_t\}}(\lambda_{X_t \cap \Lambda_{\partial B \cup \{\mathbf{u}_t\}}})}, \quad (3)$$

where the choice $C(\mathcal{U}_t, \mathbf{u}_t, X_t)$ serves three main purposes.

First, it must guarantee that the right-hand side of (3) is a probability. To achieve this we need, for $H = (\mathcal{U}_t \cup B)^c$ and almost all realizations of X_t , \mathcal{U}_t and \mathbf{u}_t , that

$$C(\mathcal{U}_t, \mathbf{u}_t, X_t) \leq \inf_{\xi \in \mathcal{N}_{\Lambda_H}} \frac{Z_{\Lambda_B \setminus \{\mathbf{u}_t\}}(\lambda_{\xi \cup (X_t \cap \Lambda_{\mathcal{U}_t})})}{Z_{\Lambda_B}(\lambda_{\xi \cup (X_t \cap \Lambda_{\mathcal{U}_t \setminus \{\mathbf{u}_t\}})})}. \quad (4)$$

Second, $C(\mathcal{U}_t, \mathbf{u}_t, X_t)$ must introduce no new bias. This is guaranteed if $C(\mathcal{U}_t, \mathbf{u}_t, X_t)$ only depends on $X_t \cap \Lambda_{\mathcal{U}_t}$. Finally, it must ensure that the algorithm terminates almost surely. It suffices to ensure $C(\mathcal{U}_t, \mathbf{u}_t, X_t)$ is uniformly bounded away from 0 for almost all realizations of X_t , implying that the same holds for the right-hand side of (3). We refer to a function $C(\cdot)$ satisfying these requirements as a *Bayes filter correction*.

If we use a Bayes filter as given in (3), keeping X_t and \mathcal{U}_t unchanged whenever $F_t = 0$ introduces new bias. To prevent this, we set $\mathcal{U}_{t+1} = \mathcal{U}_t \cup \partial B$ in step 3.b). Since the algorithm only terminates once $\mathcal{U}_t = \emptyset$, we further require the Bayes filter correction to ensure that the probability of $F_t = 0$ is small to guarantee efficiency.

Constructing a Bayes filter correction that satisfies the requirements above and allows for efficient sampling of F_t is a non-trivial task. In the next subsections, we present two approaches for this, the first specialized to the hard-sphere model without requirements, and the second one for more general potentials with strong spatial mixing of the point process.

Crucially, assuming strong spatial mixing, both constructions allow us to control the success probability of the Bayes filter via the update radius ℓ in the construction of the updated set of boxes B (see step 1).

2.1 Bayes filter for the hard-sphere model

To construct a Bayes filter for the hard-sphere model, we efficiently approximate the right-hand side of (4). To approximate the infimum over the uncountable set of configurations $\xi \in \mathcal{N}_{\Lambda_H}$ we take the minimum over a finite, but sufficiently rich set of configurations, balancing the quality of approximation with the computation required. In fact the number of configurations needed will depend only on the volume of $\Lambda_{B \cup \partial B}$. We approximate the fraction of partition functions in (4) with running time only depending on the volume of $\Lambda_{B \cup \partial B}$. As a result, we efficiently compute a Bayes filter correction $C_\varepsilon(\cdot)$, with the parameter $\varepsilon > 0$ controlling how much $C_\varepsilon(\mathcal{U}_t, \mathbf{u}_t, X_t)$ deviates from the right-hand side of (4).

While our construction of $C_\varepsilon(\cdot)$ guarantees correctness of the sampling algorithm for any $\varepsilon > 0$, proving efficiency requires more. With strong spatial mixing, we choose ε so that the probability that $F_t = 0$ is uniformly bounded above, ensuring $O(|\Lambda|)$ iterations in expectation.

It remains to argue that we can efficiently sample F_t , using the Bayes filter correction $C_\varepsilon(\cdot)$. Explicitly computing the success probability of F_t as in (3) would require computing the fraction of partition functions on the right-hand side exactly, while approximating these partition functions would require that the approximation error only depends on $X_t \cap \Lambda_{\mathcal{U}_t}$, to avoid new bias.

It is unclear how to implement these approaches, so instead we use Bernoulli factories to sample F_t without knowing the success probability. To do so, we observe that the fraction of partition functions can be written as a ratio of probabilities for drawing the empty set from a conditional hard-sphere model on Λ_B and $\Lambda_{B \setminus \{\mathbf{u}_t\}}$. Since both regions have constant volume, rejection sampling obtains Bernoulli random variables with these success probabilities in constant time, and hence construct a Bernoulli factory for F_t with constant expected running time. Wald's identity yields a total expected running time $O(|\Lambda|)$ for the algorithm.

2.2 Bayes filter for general potentials

We now consider the case of general bounded-range, repulsive potentials. Unlike the hard sphere model, it is not clear here how to approximate the infimum in (4) from a finite set of boundary configurations. However, given constants $a, b > 0$ such that ϕ satisfies (a, b) -strong spatial mixing, we can explicitly compute a function $\delta(a, b)$ so that

$$C_{a,b}(\mathcal{U}_t, \mathbf{u}_t, X_t) = \delta(a, b) \cdot \frac{Z_{\Lambda_{B \setminus \{\mathbf{u}_t\}}}(\lambda_{X_t \cap \Lambda_{\mathcal{U}_t}})}{Z_{\Lambda_B}(\lambda_{X_t \cap \Lambda_{\mathcal{U}_t \setminus \{\mathbf{u}_t\}}})}$$

is a Bayes filter correction. With strong spatial mixing, we use $C_{a,b}(\cdot)$ to construct a Bayes filter such that probability that $F_t = 0$ is bounded above, again implying a bound of $O(|\Lambda|)$ on the expected number of iterations of the algorithm.

Note that in this setting, we require spatial mixing for both correctness and efficiency, while for the hard-sphere model we only need it for efficiency. Another crucial difference is that, while we can explicitly compute $\delta(a, b)$, the same does not hold for $C_{a,b}(\cdot)$ due to the fraction of partition functions involved. Again we circumvent this by rewriting the success probability of the Bayes filter in a suitable way and applying a Bernoulli factory for sampling F_t . Finally, we point out that we do not obtain a constant bound for the expected running time of each iteration, but instead the bound depends on the number of points in $X_t \cap \Lambda_{\partial B}$. Possible dependencies between the configuration X_t and the number of iterations prevent us from bounding the total expected running time using Wald's identity. Instead, we provide tail bounds on the number of iterations and the running time of each iteration, allowing us to derive an expected total running time that is linear in the volume of Λ up to polylogarithmic factors.

3 Preliminaries

Throughout the paper, we write \mathbb{N} for the set of strictly positive integers, and we write $\mathbb{N}_0 = \mathbb{N} \cup \{0\}$. For any $k \in \mathbb{N}$, we denote by $[k]$ the set $[1, k] \cap \mathbb{N}$.

For a bounded measurable region $\Lambda \subset \mathbb{R}^d$ and any finite point configuration $\eta \in \mathcal{N}_\Lambda$, we write $|\eta|$ for the number of points in η . Note that this notation is the same that as the one we use for the volume of a region. The particular meaning will be clear from the context. Moreover, for $k \in \mathbb{N}$, we write $\binom{\eta}{k}$ for the set $\{\eta' \subseteq \eta \mid |\eta'| = k\}$.

3.1 Gibbs point processes

We introduce some additional notation for Gibbs point processes, used in the rest of the paper. Firstly, when dealing with a tuple $(x_1, \dots, x_k) \in (\mathbb{R}^d)^k$ we frequently denote it by the corresponding bold letter \mathbf{x} . Based on this, we write $d\mathbf{x}$ for $dx_1 \dots dx_k$ and $H(\mathbf{x})$ for $H(x_1, \dots, x_k)$. Moreover, for any $k \in \mathbb{N}_0$ and $\mathbf{x} = (x_1, \dots, x_k) \in (\mathbb{R}^d)^k$ we write $\eta_{\mathbf{x}}$ for the set $\{x_1, \dots, x_k\}$, where the case $k = 0$ results in $\eta_{\mathbf{x}} = \emptyset$. Finally, for $\mathbf{x} \in \Lambda^k$ we write $\boldsymbol{\lambda}^{\mathbf{x}}$ for $\prod_{i \in [k]} \boldsymbol{\lambda}(x_i)$. This simplifies the definition of $\mu_{\boldsymbol{\lambda}}$ given in the introduction to

$$\mu_{\boldsymbol{\lambda}}(A) = \frac{1}{Z_{\Lambda}(\boldsymbol{\lambda})} \sum_{k \geq 0} \frac{1}{k!} \int_{\Lambda^k} \mathbb{1}_{\eta_{\mathbf{x}} \in A} \boldsymbol{\lambda}^{\mathbf{x}} e^{-H(\mathbf{x})} d\mathbf{x}.$$

Next, we formalize two different notions of restricting a Gibbs point process on Λ to a subregion $\Lambda' \subseteq \Lambda$ that are relevant throughout the paper.

The first is based on restricting the support of $\boldsymbol{\lambda}$ by defining a new activity function $\boldsymbol{\lambda} \mathbb{1}_{\Lambda'} : y \mapsto \boldsymbol{\lambda}(y) \cdot \mathbb{1}_{y \in \Lambda'}$ (for constant activity λ , we write $\lambda \mathbb{1}_{\Lambda'} : y \mapsto \lambda \mathbb{1}_{y \in \Lambda'}$). The resulting

Gibbs point process is a probability measure on $(\mathcal{N}_\Lambda, \mathfrak{R}_\Lambda)$ with

$$\begin{aligned}\mu_{\lambda \mathbf{1}_{\Lambda'}}(A) &= \frac{1}{Z_\Lambda(\lambda \mathbf{1}_{\Lambda'})} \sum_{k \geq 0} \frac{1}{k!} \int_{\Lambda^k} \mathbf{1}_{\eta_{\mathbf{x}} \in A} (\lambda \mathbf{1}_{\Lambda'})^{\mathbf{x}} e^{-H(\mathbf{x})} d\mathbf{x} \\ &= \frac{1}{Z_{\Lambda'}(\lambda)} \sum_{k \geq 0} \frac{1}{k!} \int_{\Lambda'^k} \mathbf{1}_{\eta_{\mathbf{x}} \in A} \lambda^{\mathbf{x}} e^{-H(\mathbf{x})} d\mathbf{x}\end{aligned}$$

for all $A \in \mathfrak{R}_\Lambda$. In particular, for $A = \{\eta \in \mathcal{N}_\Lambda \mid \eta \cap (\Lambda')^c > 0\}$, it holds that $\mu_{\lambda \mathbf{1}_{\Lambda'}}(A) = 0$.

The second way of restricting a Gibbs point process μ_λ is by projecting it to a measurable subregion $\Lambda' \subseteq \Lambda$. To this end, we write $\mu_\lambda[\Lambda']$ for the image measure of μ_λ under the map $\mathcal{N}_\Lambda \rightarrow \mathcal{N}_{\Lambda'}$, $\eta \mapsto \eta \cap \Lambda'$. By construction, $\mu_\lambda[\Lambda']$ is a probability distribution on $(\mathcal{N}_{\Lambda'}, \mathfrak{R}_{\Lambda'})$ that, for every $A \in \mathfrak{R}_{\Lambda'}$, assigns a probability

$$\mu_\lambda[\Lambda'](A) = \frac{1}{Z_\Lambda(\lambda)} \sum_{k \geq 0} \frac{1}{k!} \int_{\Lambda^k} \mathbf{1}_{\eta_{\mathbf{x}} \cap \Lambda' \in A} \lambda^{\mathbf{x}} e^{-H(\mathbf{x})} d\mathbf{x}.$$

As discussed in Section 1, we frequently modify the activity function to encode the effect of fixing a certain point set (boundary condition). More precisely, for a fixed potential ϕ , an activity function λ and a point set $\eta \in \mathcal{N}_\Lambda$ we write λ_η for the function $y \mapsto \lambda(y) e^{-\sum_{x \in \eta} \phi(x, y)}$. Similarly, for $k \in \mathbb{N}$ and $\mathbf{x} \in \Lambda^k$ we write $\lambda_{\mathbf{x}}$ for the activity function $y \mapsto \lambda(y) e^{-\sum_{i \in [k]} \phi(x_i, y)}$. We extend this notation to constant activity $\lambda \in \mathbb{R}_{\geq 0}$, writing $\lambda_\eta : y \mapsto \lambda e^{-\sum_{x \in \eta} \phi(x, y)}$ and $\lambda_{\mathbf{x}} : y \mapsto \lambda e^{-\sum_{i \in [k]} \phi(x_i, y)}$. Using this notation, a useful alternative definition of $\mu_\lambda[\Lambda']$ is given by

$$\begin{aligned}\mu_\lambda[\Lambda'](A) &= \frac{1}{Z_\Lambda(\lambda)} \sum_{k \geq 0} \frac{1}{k!} \int_{\Lambda'^k} \mathbf{1}_{\eta_{\mathbf{x}} \in A} \lambda^{\mathbf{x}} e^{-H(\mathbf{x})} Z_\Lambda(\lambda_{\mathbf{x}} \mathbf{1}_{(\Lambda')^c}) d\mathbf{x} \\ &= \frac{1}{Z_\Lambda(\lambda)} \sum_{k \geq 0} \frac{1}{k!} \int_{\Lambda'^k} \mathbf{1}_{\eta_{\mathbf{x}} \in A} \lambda^{\mathbf{x}} e^{-H(\mathbf{x})} Z_{(\Lambda')^c}(\lambda_{\mathbf{x}}) d\mathbf{x}\end{aligned}$$

for $A \in \mathfrak{R}_{\Lambda'}$. In particular, note that

$$\mu_{\lambda \mathbf{1}_{\Lambda'}}[\Lambda'](A) = \frac{1}{Z_{\Lambda'}(\lambda)} \sum_{k \geq 0} \frac{1}{k!} \int_{\Lambda'^k} \mathbf{1}_{\eta_{\mathbf{x}} \in A} \lambda^{\mathbf{x}} e^{-H(\mathbf{x})} d\mathbf{x}.$$

While $\mu_{\lambda \mathbf{1}_{\Lambda'}}[\Lambda']$ and $\mu_{\lambda \mathbf{1}_{\Lambda'}}$ seem similar, the difference is that the former is a distribution on $(\mathcal{N}_{\Lambda'}, \mathfrak{R}_{\Lambda'})$ whereas the latter is defined on $(\mathcal{N}_\Lambda, \mathfrak{R}_\Lambda)$.

We introduce further concepts related to Gibbs point processes, such a point density functions when they are required. Moreover, various useful properties of Gibbs point processes are given in Appendix C.

3.2 Bernoulli factories

In designing our sampling algorithm, it will be useful to consider the following Bernoulli factory problem. We are given access to a sampler for $\text{Ber}(p)$ and for $\text{Ber}(q)$, that is samplers of Bernoulli random variables with parameters p and q respectively, where we further assume $p < q$. We want to sample a random variable $Z \sim \text{Ber}\left(\frac{p}{q}\right)$.

Most work on Bernoulli factories studies their running time in terms of the number of coin flips required. In our setting, the time needed to generate each of these coin flips is random variable. Fortunately, suitable independence assumptions hold in our setting allowing us to prove the following lemma.

Lemma 3.1. *Fix some $p, q \in [0, 1]$ such that $q - p \geq \epsilon$ for some $\epsilon > 0$. Further assume that we have oracle access to a sampler from $\text{Ber}(p)$ and $\text{Ber}(q)$ in the following sense:*

1. *every sample from $\text{Ber}(p)$ (resp. $\text{Ber}(q)$) is independent from all previous samples;*
2. *the expected running time for obtaining a sample from $\text{Ber}(p)$ (resp. $\text{Ber}(q)$), conditioned on previously obtained samples, is uniformly bounded by some $t \in \mathbb{R}_{\geq 0}$.*

Then we can sample from $\text{Ber}\left(\frac{p}{q}\right)$ in $O(t\epsilon^{-2})$ expected time.

Lemma 3.1, will play a key role in bounding the expected running time of our algorithm. To proceed with the formal description of our algorithm, we defer the proof of this lemma to Section 9.

4 The algorithm

Let $\Lambda = [0, L]^d$ and consider a Gibbs point processes on Λ with uniform activity $\lambda(x) \equiv \lambda$ for some $\lambda \in \mathbb{R}_{>0}$ and repulsive potential ϕ with finite range $r \in \mathbb{R}_{>0}$. Throughout the analysis of our algorithm, it will be useful to focus on configurations $\eta \in \mathcal{N}_\Lambda$ such that $\phi(x, y) < \infty$ for all $\{x, y\} \in \binom{\eta}{2}$, in which case we call η a *feasible configuration*.

Before stating our algorithm, we first formalize how we divide Λ into smaller boxes, following the description given in Section 2. For a r and L as above, let $N = \lceil L/r \rceil$. We set $\mathcal{V} = \{0, \dots, N-1\}^d$ to be the set of box indices and associate each box index $\mathbf{v} = (v_1, \dots, v_d) \in \mathcal{V}$ with the region $\Lambda_{\mathbf{v}} = ([v_1 r, (v_1 + 1)r) \times \dots \times [v_d r, (v_d + 1)r) \cap \Lambda$. As in Section 2, we extend this notation to sets of box indices $S \subseteq \mathcal{V}$ by setting $\Lambda_S = \bigcup_{\mathbf{v} \in S} \Lambda_{\mathbf{v}}$. Further, recall that, for $\mathbf{v} \in \mathcal{V}$, we write $\mathbb{B}_k(\mathbf{v})$ for the set of boxes $\mathbf{w} \in \mathcal{V}$ with $\|\mathbf{v} - \mathbf{w}\|_\infty < k$. As mentioned earlier, our algorithm tries to update in each step the point configuration on a subset of boxes $B \subseteq \mathcal{V}$. To this end, for $S \subseteq \mathcal{V}$, $\mathbf{v} \in S$, $r \in \mathbb{R}_{>0}$ and $\ell \in \mathbb{N}$, we define

$$B(S, \mathbf{v}, \ell) := \{\mathbf{v}\} \cup (\mathbb{B}_\ell(\mathbf{v}) \cap S).$$

We refer to the parameter ℓ as the *update radius*. Finally, recall that we write $\partial S = (\bigcup_{\mathbf{v} \in S} \mathbb{B}_1(\mathbf{v})) \setminus S$ for the outer boundary of $S \subseteq \mathcal{V}$.

If the algorithm updates the point configuration in any iteration t depends on the outcome of a Bernoulli random variable F_t , called the *Bayes filter*. We introduce the following definition.

Definition 4.1. Fix a repulsive potential ϕ of range $r \in \mathbb{R}_{>0}$, an activity $\lambda \in \mathbb{R}_{>0}$ and some $\ell \in \mathbb{N}$. We call a function $C : 2^{\mathcal{V}} \times \mathcal{V} \times \mathcal{N}_\Lambda \rightarrow [0, 1]$ a Bayes filter correction if, for all non-empty $S \subseteq \mathcal{V}$ and $\mathbf{v} \in S$, it holds that

1. $C(S, \mathbf{v}, \cdot)$ is \mathfrak{R}_{Λ_S} -measurable (in particular $C(S, \mathbf{v}, \eta) = C(S, \mathbf{v}, \eta \cap \Lambda_S)$ for all $\eta \in \mathcal{N}_\Lambda$),
2. there is some $\varepsilon > 0$ such that for $B = B(S, \mathbf{v}, \ell)$, $H = (S \cup B)^c$ and all feasible $\eta \in \mathcal{N}_\Lambda$ it holds that

$$\varepsilon \leq C(S, \mathbf{v}, \eta) \leq \inf_{\substack{\xi \in \mathcal{N}_{\Lambda_H} \\ \xi \cup (\eta \cap \Lambda_S) \text{ is feasible}}} \left\{ \frac{Z_{\Lambda_B \setminus \{\mathbf{v}\}}(\lambda_{\xi \cup (\eta \cap \Lambda_S)})}{Z_{\Lambda_B}(\lambda_{\xi \cup (\eta \cap \Lambda_S \setminus \{\mathbf{v}\})})} \right\}.$$

Our perfect sampling procedure is stated in Algorithm 1.

Algorithm 1: Perfect sampling algorithm for repulsive Gibbs point processes

Data: region $\Lambda = [0, L]^d$, repulsive potential ϕ of range at most $r \in \mathbb{R}_{>0}$, activity $\lambda \in \mathbb{R}_{>0}$, update radius $\ell \in \mathbb{N}$

- 1 set $t = 0$, $\mathcal{U}_t = \mathcal{V}$, $X_t = \emptyset$
- 2 **while** $\mathcal{U}_t \neq \emptyset$ **do**
- 3 draw $\mathbf{u}_t \in \mathcal{U}_t$ uniformly at random
- 4 set $B = B(\mathcal{U}_t, \mathbf{u}_t, \ell)$
- 5 draw F_t from $\text{Ber} \left(C(\mathcal{U}_t, \mathbf{u}_t, X_t) \cdot \frac{Z_{\Lambda_B}(\lambda_{X_t \cap \Lambda_{\partial B}})}{Z_{\Lambda_B \setminus \{\mathbf{u}_t\}}(\lambda_{X_t \cap \Lambda_{\partial B} \cup \{\mathbf{u}_t\}})} \right)$ where C is a Bayes filter correction as in Definition 4.1
- 6 **if** $F_t = 1$ **then**
- 7 draw Y from $\mu_{\lambda_{X_t \cap (\Lambda_B)^c}} \mathbb{1}_{\Lambda_B}[\Lambda_B]$
- 8 set $X_{t+1} = (X_t \setminus \Lambda_B) \cup Y$
- 9 set $\mathcal{U}_{t+1} = \mathcal{U}_t \setminus \{\mathbf{u}_t\}$
- 10 **else**
- 11 set $\mathcal{U}_{t+1} = \mathcal{U}_t \cup \partial B$
- 12 increase t by 1
- 13 **return** X_t

Before we analyze Algorithm 1, we first argue that the success probability of the Bayes filter that we require in line 5 is indeed a well defined probability. For this, we use the following lemma.

Lemma 4.2. *Suppose C is a Bayes filter correction. Let $S \subseteq \mathcal{V}$ be non-empty, $\mathbf{v} \in S$ and $B = B(S, \mathbf{v}, \ell)$. There is some $\varepsilon > 0$ such that, for all feasible $\eta \in \mathcal{N}_\Lambda$, it holds that*

$$\varepsilon \leq C(S, \mathbf{v}, \eta) \cdot \frac{Z_{\Lambda_B}(\lambda_{\eta \cap \Lambda_{\partial B}})}{Z_{\Lambda_{B \setminus \{\mathbf{v}\}}}(\lambda_{\eta \cap \Lambda_{\partial B \cup \{\mathbf{v}\}}})} \leq 1.$$

Proof. Fix S and \mathbf{v} . For the lower bound, note that for all $\eta \in \mathcal{N}_\Lambda$

$$Z_{\Lambda_B}(\lambda_{\eta \cap \Lambda_{\partial B}}) \geq Z_{\Lambda_{B \setminus \{\mathbf{v}\}}}(\lambda_{\eta \cap \Lambda_{\partial B}}) \geq Z_{\Lambda_{B \setminus \{\mathbf{v}\}}}(\lambda_{\eta \cap \Lambda_{\partial B \cup \{\mathbf{v}\}}}).$$

Thus, by the definition of a Bayes filter correction, there is some $\varepsilon > 0$ such that for all feasible $\eta \in \mathcal{N}_\Lambda$

$$C(S, \mathbf{v}, \eta) \cdot \frac{Z_{\Lambda_B}(\lambda_{\eta \cap \Lambda_{\partial B}})}{Z_{\Lambda_{B \setminus \{\mathbf{v}\}}}(\lambda_{\eta \cap \Lambda_{\partial B \cup \{\mathbf{v}\}}})} \geq C(S, \mathbf{v}, \eta) \geq \varepsilon.$$

To derive the upper bound, note that by Lemma C.2 $Z_{\Lambda_B}(\lambda_{\eta \cap \Lambda_{\partial B}}) = Z_{\Lambda_B}(\lambda_{\eta \cap (\Lambda_B)^c})$ and $Z_{\Lambda_{B \setminus \{\mathbf{v}\}}}(\lambda_{\eta \cap \Lambda_{\partial B \cup \{\mathbf{v}\}}}) = Z_{\Lambda_{B \setminus \{\mathbf{v}\}}}(\lambda_{\eta \cap (\Lambda_{B \setminus \{\mathbf{v}\}})^c})$. Note that, for feasible $\eta \in \mathcal{N}_\Lambda$, it further holds that $(\eta \cap \Lambda_H) \cup (\eta \cap \Lambda_S)$ is feasible as well. By the definition of a Bayes filter correction, this implies

$$C(S, \mathbf{v}, \eta) \leq \frac{Z_{\Lambda_{B \setminus \{\mathbf{v}\}}}(\lambda_{\eta \cap (\Lambda_{B \setminus \{\mathbf{v}\}})^c})}{Z_{\Lambda_B}(\lambda_{\eta \cap (\Lambda_B)^c})}.$$

Consequently, it holds that

$$C(S, \mathbf{v}, \eta) \cdot \frac{Z_{\Lambda_B}(\lambda_{\eta \cap \Lambda_{\partial B}})}{Z_{\Lambda_{B \setminus \{\mathbf{v}\}}}(\lambda_{\eta \cap \Lambda_{\partial B \cup \{\mathbf{v}\}}})} = C(S, \mathbf{v}, \eta) \cdot \frac{Z_{\Lambda_B}(\lambda_{\eta \cap (\Lambda_B)^c})}{Z_{\Lambda_{B \setminus \{\mathbf{v}\}}}(\lambda_{\eta \cap (\Lambda_{B \setminus \{\mathbf{v}\}})^c})} \leq 1,$$

which proves the claim □

We use the previous lemma to derive the following statement.

Lemma 4.3. *For all $t \in \mathbb{N}_0$ the following holds almost surely*

- 1) X_t is a feasible configuration
- 2) there is some $\varepsilon > 0$ such that, given $\mathcal{U}_t \neq \emptyset$,

$$\varepsilon \leq C(\mathcal{U}_t, \mathbf{u}_t, X_t) \cdot \frac{Z_{\Lambda_B}(\lambda_{X_t \cap \Lambda_{\partial B}})}{Z_{\Lambda_{B \setminus \{\mathbf{u}_t\}}}(\lambda_{X_t \cap \Lambda_{\partial B \cup \{\mathbf{u}_t\}}})} \leq 1.$$

Proof. We prove this statement via induction over the iteration $t \in \mathbb{N}_0$. For $t = 0$, note that $X_0 = \emptyset$. Thus, 1) is trivially true. Moreover, 2) follows from applying Lemma 4.2 to $C(\mathcal{V}, \mathbf{v}, \emptyset)$ for every $\mathbf{v} \in \mathcal{V}$.

Now, suppose our claim holds at some iteration $t \in N$. We start by showing that 1) holds in iteration $t + 1$. First, note that if $\mathcal{U}_t = \emptyset$, then there is nothing to prove. Thus, we may assume $\mathcal{U}_t \neq \emptyset$. If $F_t = 0$, then $X_{t+1} = X_t$. Thus, in this case, X_{t+1} is feasible if and only if X_t was feasible, which holds almost surely by the induction hypothesis. Next, consider the case $F_t = 1$ and set $B = B(\mathcal{U}_t, \mathbf{u}_t, \ell)$. Note that

$$\sum_{\{x,y\} \in \binom{X_{t+1}}{2}} \phi(x,y) = \sum_{\{x,y\} \in \binom{X_{t+1} \cap (\Lambda_B)^c}{2}} \phi(x,y) + \sum_{\{x,y\} \in \binom{X_{t+1} \cap \Lambda_B}{2}} \phi(x,y) + \sum_{\substack{x \in X_{t+1} \cap \Lambda_B \\ y \in X_{t+1} \cap (\Lambda_B)^c}} \phi(x,y).$$

Since $X_{t+1} \cap (\Lambda_B)^c = X_t \cap (\Lambda_B)^c$ and X_t is almost surely feasible, we have

$$\sum_{\{x,y\} \in \binom{X_{t+1} \cap (\Lambda_B)^c}{2}} \phi(x,y) < \infty$$

with probability 1. Moreover, note that $X_{t+1} \cap \Lambda_B \sim \mu_{\lambda_{X_t \cap (\Lambda_B)^c} \mathbb{1}_{\Lambda_B}}[\Lambda_B]$. Thus, it also holds that

$$\begin{aligned} \sum_{\{x,y\} \in \binom{X_{t+1} \cap \Lambda_B}{2}} \phi(x,y) + \sum_{\substack{x \in X_{t+1} \cap \Lambda_B \\ y \in X_{t+1} \cap (\Lambda_B)^c}} \phi(x,y) &= \sum_{\{x,y\} \in \binom{X_{t+1} \cap \Lambda_B}{2}} \phi(x,y) + \sum_{\substack{x \in X_{t+1} \cap \Lambda_B \\ y \in X_t \cap (\Lambda_B)^c}} \phi(x,y) \\ &< \infty \end{aligned}$$

with probability 1. Consequently, X_{t+1} is almost surely feasible, proving 1). For 2), assume that $\mathcal{U}_{t+1} \neq \emptyset$. Applying Lemma 4.2 for every non-empty $S \subseteq \mathcal{V}$ and $\mathbf{v} \in S$ yields the desired bounds on $C(\mathcal{U}_{t+1}, \mathbf{u}_{t+1}, X_{t+1})$ whenever X_{t+1} is feasible. As we have just shown, this is the case almost surely, which concludes the proof. \square

Considering Algorithm 1, an immediate question is how to construct the Bayes filter correction in line 5, and in particular, how to do so in such a way that the Bayes filter F_t can be sampled efficiently. However, we will defer this question for now and first prove that Algorithm 1 produces the correct output distribution.

5 Proof of correctness

In this section we prove that Algorithm 1 produces the correct output distribution. That is for $T = \inf\{t \in \mathbb{N}_0 \mid \mathcal{U}_t = \emptyset\}$ it holds that $X_T \sim \mu_\lambda$. We first show that the number of iterations T is finite almost surely. This directly follows as a corollary of Lemma 4.3.

Corollary 5.1. *Algorithm 1 terminates almost surely after finitely many iterations. That is, for $T = \inf\{t \in \mathbb{N}_0 \mid \mathcal{U}_t = \emptyset\}$ we have $\mathbb{P}[T < \infty] = 1$.*

Proof. By Lemma 4.3, the probability that $|\mathcal{U}_t|$ decreases in each step is uniformly bounded away from 0. Thus, there is a positive probability of going from any \mathcal{U}_t to the empty set in $|\mathcal{U}_t|$ steps. Since $|\mathcal{U}_t| \leq |\mathcal{V}|$, this means for every $k \in \mathbb{N}$, it holds that the probability that $\mathcal{U}_{k \cdot |\mathcal{V}|} = \emptyset$ is bounded away from 0 uniformly in k . Thus, T is dominated by a geometric random variable with strictly positive success probability, which proves the claim. \square

To analyze Algorithm 1, we consider it as a Markov chain $(X_t, \mathcal{U}_t, F_t, \mathbf{u}_t)_{t \in \mathbb{N}_0}$, which we set to remain constant for $t \geq T$. Throughout this section, we write $\Omega = (\mathcal{N}_\Lambda \times 2^\mathcal{V} \times \{0, 1\} \times \mathcal{V})^{\mathbb{N}_0}$ for the state space of all trajectories of that Markov chain. Further, we equip Ω with the σ -field $\mathcal{A} = (\mathfrak{R}_\Lambda \otimes 2^{2^\mathcal{V}} \otimes 2^{\{0,1\}} \otimes 2^\mathcal{V})^{\otimes \mathbb{N}_0}$, and we denote by \mathbb{P} the distribution on (Ω, \mathcal{A}) induced by Algorithm 1.

Before going into the technical part of proving correctness, a few remarks about our notation are in place. Firstly, for any $A \in \mathcal{A}$ with $\mathbb{P}[A] > 0$ we write \mathbb{P}_A as shorthand for the probability measure $\mathbb{P}_A[\cdot] = \mathbb{P}[\cdot | A]$. Note that for all events $A, B \in \mathcal{A}$ with $\mathbb{P}[A \cap B] > 0$ it holds that $\mathbb{P}_{A \cap B} = (\mathbb{P}_A)_B$. Throughout our proofs, we use conditional expectations to make conditioning on partial point configurations $X_t \cap \Lambda'$ rigorous. In particular, we frequently condition on a sub- σ -field $\mathcal{F} \subseteq \mathcal{A}$ and an event $A \in \mathcal{A}$ with $\mathbb{P}[A] > 0$ at the same time. Formally, for a measurable function $f : \Omega \rightarrow \mathbb{R}$ we write $\mathbb{E}[f | \mathcal{F}; A]$ for the conditional expectation of f given \mathcal{F} under the conditional measure \mathbb{P}_A . Note that any identity involving $\mathbb{E}[f | \mathcal{F}; A]$ should be understood to hold \mathbb{P}_A -almost surely. More details on conditional expectations can be found in Appendix A.1. Moreover, we often make use of the concept of regular condition distributions. For more details, see Appendix A.2. Lastly, for every bounded measurable $\Lambda'_1 \subseteq \Lambda'_2 \subseteq \mathbb{R}^d$, we write $\pi_{\Lambda'_1, \Lambda'_2} : \mathcal{N}_{\Lambda'_1} \rightarrow \mathcal{N}_{\Lambda'_2}$ for the projection $\eta \mapsto \eta \cap \Lambda'_2$.

Our main result in this section is the following statement.

Theorem 5.2. *For all $t \in \mathbb{N}_0$ with $\mathbb{P}[\mathcal{U}_t = \emptyset] > 0$ and all $A \in \mathfrak{R}_\Lambda$, it holds that*

$$\mathbb{P}[X_t \in A | \mathcal{U}_t = \emptyset] = \mu_\lambda(A).$$

Since the algorithm terminates when $\mathcal{U}_t = \emptyset$, this implies that the output of Algorithm 1 follows the distribution $\mu_\lambda(A)$. We deduce Theorem 5.2 from the following invariant.

Lemma 5.3. *For all $t \in \mathbb{N}_0$, $S \subseteq \mathcal{V}$ with $\mathbb{P}[\mathcal{U}_t = S] > 0$ and $A \in \mathfrak{R}_{(\Lambda_S)^c}$ it holds that*

$$\mathbb{E}[\mathbf{1}_{X_t \cap (\Lambda_S)^c \in A} | X_t \cap \Lambda_S; \mathcal{U}_t = S] = \mu_{\lambda_{X_t \cap \Lambda_S}}[\mathbf{1}_{(\Lambda_S)^c}](A).$$

In particular, the map

$$(\omega, A) \mapsto \mu_{\lambda_{X_t(\omega) \cap \Lambda_S}}[\mathbf{1}_{(\Lambda_S)^c}](A) \quad \omega \in \Omega, A \in \mathfrak{R}_{(\Lambda_S)^c}$$

is a regular conditional distribution of $X_t \cap (\Lambda_S)^c$ given $\sigma(X_t \cap \Lambda_S)$ under the probability measure $\mathbb{P}_{\{\mathcal{U}_t = S\}}$.

Before we get into proving Lemma 5.3, we first show how Theorem 5.2 follows from it.

Proof of Theorem 5.2. Note that $\Lambda_\emptyset = \emptyset$. Thus, it holds that $(\Lambda_\emptyset)^c = \Lambda$ and $\sigma(X_t \cap \Lambda_\emptyset) = \{\emptyset, \Omega\}$ for all $t \in \mathbb{N}_0$. Using Lemma 5.3, we obtain for all $A \in \mathfrak{R}_\Lambda$

$$\begin{aligned} \mathbb{P}[X_t \in A \mid \mathcal{U}_t = \emptyset] &= \mathbb{E}[\mathbb{1}_{X_t \in A} \mid \mathcal{U}_t = \emptyset] \\ &= \mathbb{E}[\mathbb{1}_{X_t \in A} \mid X_t \cap \Lambda_\emptyset; \mathcal{U}_t = \emptyset] \\ &= \mu_\lambda(A), \end{aligned}$$

which proves the theorem. \square

We proceed by stating and proving several lemmas that we will use to prove Lemma 5.3.

Lemma 5.4. *Fix $t \in \mathbb{N}_0$, and assume that for all $S \subseteq \mathcal{V}$ with $\mathbb{P}[\mathcal{U}_t = S] > 0$ and all $A \in \mathfrak{R}_{(\Lambda_S)^c}$ it holds that*

$$\mathbb{E}[\mathbb{1}_{X_t \cap (\Lambda_S)^c \in A} \mid X_t \cap \Lambda_S; \mathcal{U}_t = S] = \mu_{\lambda_{X_t \cap \Lambda_S} \mathbb{1}_{(\Lambda_S)^c}}[(\Lambda_S)^c](A).$$

Let $E = \{\mathcal{U}_t = S, \mathbf{u}_t = \mathbf{v}\}$ for some $S \in 2^\mathcal{V} \setminus \{\emptyset\}$ and $\mathbf{v} \in S$ such that $\mathbb{P}[E] > 0$. For any measurable region $\Lambda' \subseteq (\Lambda_S)^c$ and any event $A \in \mathfrak{R}_{\Lambda'}$ it holds that

$$\mathbb{E}[\mathbb{1}_{X_t \cap \Lambda' \in A} \mid X_t \cap \Lambda_S; E] = \mu_{\lambda_{X_t \cap \Lambda_S} \mathbb{1}_{(\Lambda_S)^c}}[\Lambda'](A).$$

In particular,

$$(\omega, A) \mapsto \mu_{\lambda_{X_t \cap \Lambda_S} \mathbb{1}_{(\Lambda_S)^c}}[\Lambda'](A) \quad \omega \in \Omega, A \in \mathfrak{R}_{\Lambda'}$$

is a regular conditional distribution of $X_t \cap \Lambda'$ given $\sigma(X_t \cap \Lambda_S)$ under \mathbb{P}_E .

Proof. Fix some measurable region $\Lambda' \subseteq (\Lambda_S)^c$ and note that $\mu_{\lambda_{X_t \cap \Lambda_S} \mathbb{1}_{(\Lambda_S)^c}}[\Lambda']$ is a probability distribution on $(\mathcal{N}_{\Lambda'}, \mathfrak{R}_{\Lambda'})$. Thus, it suffices to show that $\mu_{\lambda_{X_t \cap \Lambda_S} \mathbb{1}_{(\Lambda_S)^c}}[\Lambda'](A)$ is also a version of the conditional expectation $\mathbb{E}[\mathbb{1}_{X_t \cap (\Lambda')^c \in A} \mid X_t \cap \Lambda_S; E]$ for all events $A \in \mathfrak{R}_{\Lambda'}$. Write $\pi_{\Lambda'}$ as shorthand for the projection $\pi_{(\Lambda_S)^c, \Lambda'}$. By the assumptions of the lemma, we have

$$\begin{aligned} \mathbb{E}[\mathbb{1}_{X_t \cap \Lambda' \in A} \mid X_t \cap \Lambda_S; \mathcal{U}_t = S] &= \mathbb{E}\left[\mathbb{1}_{X_t \cap (\Lambda_S)^c \in \pi_{\Lambda'}^{-1}(A)} \mid X_t \cap \Lambda_S; \mathcal{U}_t = S\right] \\ &= \mu_{\lambda_{X_t \cap \Lambda_S} \mathbb{1}_{(\Lambda_S)^c}}[(\Lambda_S)^c](\pi_{\Lambda'}^{-1}(A)) \\ &= \mu_{\lambda_{X_t \cap \Lambda_S} \mathbb{1}_{(\Lambda_S)^c}}[\Lambda'](A). \end{aligned}$$

Next, we use Lemma A.3 part (2) to argue that this still holds if we change the probability measure from $\mathbb{P}_{\{\mathcal{U}_t = S\}}$ to \mathbb{P}_E . Note that, given $\mathcal{U}_t = S$, \mathbf{u}_t is chosen uniformly from S independent of X_t . Therefore, we have

$$\begin{aligned} \mathbb{E}[\mathbb{1}_{\mathbf{u}_t = \mathbf{v}} \mid X_t \cap \Lambda_S; \mathcal{U}_t = S] &= \mathbb{E}[\mathbb{1}_{\mathbf{u}_t = \mathbf{v}} \mid \mathcal{U}_t = S] \\ &= \mathbb{E}[\mathbb{1}_{\mathbf{u}_t = \mathbf{v}} \mid X_t \cap \Lambda_S, X_t \cap \Lambda'; \mathcal{U}_t = S]. \end{aligned}$$

\square

Lemma 5.5. *Under the assumptions of Lemma 5.4, let $E_1 = \{\mathcal{U}_t = S, \mathbf{u}_t = \mathbf{v}, F_t = 1\}$ for some $S \in 2^\mathcal{V} \setminus \{\emptyset\}$ and $\mathbf{v} \in S$ such that $\mathbb{P}[E_1] > 0$, and set $R = S \setminus \{\mathbf{v}\}$. For all $A \in \mathfrak{R}_{(\Lambda_R)^c}$ it holds that*

$$\mathbb{E}[\mathbf{1}_{X_{t+1} \cap (\Lambda_R)^c \in A} \mid X_{t+1} \cap \Lambda_R; E_1] = \mu_{\lambda_{X_{t+1} \cap \Lambda_R} \mathbf{1}_{(\Lambda_R)^c}}[(\Lambda_R)^c](A).$$

Proof. Let $E = \{\mathcal{U}_t = S, \mathbf{u}_t = \mathbf{v}\}$ for S and \mathbf{v} as in the definition of E_1 , and set $B = B(S, \mathbf{v}, \ell)$ and $H = \mathcal{V} \setminus (S \cup B)$. The main idea of the proof is to prove the claim for pairs of events $A_B \in \mathfrak{R}_B, A_H \in \mathfrak{R}_H$ and use the fact that such events form a generating π -system since $(R)^c = \mathcal{U} \cup H$. To establish this, we derive a suitable version of the conditional expectation $\mathbb{E}[\mathbf{1}_{X_t \cap \Lambda_H \in A_H} \cdot \mathbf{1}_{X_{t+1} \cap \Lambda_B \in A_B} \cdot F_t \mid X_t \cap \Lambda_S; E]$, which we do by first deriving an expression for $\mathbb{E}[\mathbf{1}_{X_{t+1} \cap \Lambda_B \in A_B} \cdot F_t \mid X_t \cap \Lambda_S; E]$ and integrating it over A_H against a regular conditional distribution for $X_t \cap \Lambda_H$ given $\sigma(X_t \cap \Lambda_S)$ under the measure \mathbb{P}_E .

We start by deriving an expression for $\mathbb{E}[\mathbf{1}_{X_{t+1} \cap \Lambda_B \in A_B} \cdot F_t \mid X_t \cap \Lambda_S; E]$. Note that given the event E_1 , the point set $X_{t+1} \cap \Lambda_B$ is sampled from $\mu_{\lambda_{X_t \cap (\Lambda_B)^c} \mathbf{1}_{\Lambda_B}}[\Lambda_B]$. Therefore we have for all $A_B \in \mathfrak{R}_{\Lambda_B}$ that

$$\mathbb{E}[\mathbf{1}_{X_{t+1} \cap \Lambda_B \in A_B} \mid X_t \cap (\Lambda_B)^c; E_1] = \mu_{\lambda_{X_t \cap (\Lambda_B)^c} \mathbf{1}_{\Lambda_B}}[\Lambda_B].$$

Since the above is true independently of the partial configuration $X_t \cap \Lambda_{\mathbf{v}}$, we may add it to the condition and obtain

$$\begin{aligned} \mathbb{E}[\mathbf{1}_{X_{t+1} \cap \Lambda_B \in A_B} \mid X_t \cap (\Lambda_B \setminus \{\mathbf{v}\})^c; E_1] &= \mathbb{E}[\mathbf{1}_{X_{t+1} \cap \Lambda_B \in A_B} \mid X_t \cap (\Lambda_B)^c, X_t \cap \Lambda_{\mathbf{v}}; E_1] \\ &= \mu_{\lambda_{X_t \cap (\Lambda_B)^c} \mathbf{1}_{\Lambda_B}}[\Lambda_B](A_B) \end{aligned}$$

Note that the right-hand side of the latter equality is a function on Ω , and denote this function by g . Moreover, given E and $X_t \cap (\Lambda_B)^c$, F_t is drawn as a Bernoulli random variable with success probability

$$C(S, \mathbf{v}, X_t) \cdot \frac{Z_{\Lambda_B}(\lambda_{X_t \cap \Lambda_{\partial B}})}{Z_{\Lambda_B \setminus \{\mathbf{v}\}}(\lambda_{X_t \cap \Lambda_{\partial B \cup \{\mathbf{v}\}})},$$

where C is a Bayes filter correction as in Definition 4.1. Thus, it holds that

$$\begin{aligned} \mathbb{E}[F_t \mid X_t \cap (\Lambda_B \setminus \{\mathbf{v}\})^c; E] &= C(S, \mathbf{v}, X_t) \cdot \frac{Z_{\Lambda_B}(\lambda_{X_t \cap \Lambda_{\partial B}})}{Z_{\Lambda_B \setminus \{\mathbf{v}\}}(\lambda_{X_t \cap \Lambda_{\partial B \cup \{\mathbf{v}\}})} \\ &= C(S, \mathbf{v}, X_t \cap \Lambda_S) \cdot \frac{Z_{\Lambda_B}(\lambda_{X_t \cap \Lambda_{\partial B}})}{Z_{\Lambda_B \setminus \{\mathbf{v}\}}(\lambda_{X_t \cap \Lambda_{\partial B \cup \{\mathbf{v}\}})}, \end{aligned}$$

where the second equality holds since $C(S, \mathbf{v}, \cdot) = C(S, \mathbf{v}, \cdot \cap \Lambda_S)$ by definition. We denote the function on the right-hand side of the last equality by h . By Lemma A.1, we have

$$\mathbb{E}[\mathbf{1}_{X_{t+1} \cap \Lambda_B \in A_B} \cdot F_t \mid X_t \cap (\Lambda_B \setminus \{\mathbf{v}\})^c; E](\omega) = g(\omega) \cdot h(\omega)$$

for \mathbb{P}_E -almost all $\omega \in \Omega$.

Our next step is to use Theorem A.4 to derive an expression for the conditional expectation $\mathbb{E}[\mathbb{1}_{X_t \cap \Lambda_H \in A_H} \cdot \mathbb{1}_{X_{t+1} \cap \Lambda_B \in A_B} \cdot F_t \mid X_t \cap \Lambda_S; E]$ for every $A_H \in \mathfrak{A}_{\Lambda_H}$. To this end, note that

$$\begin{aligned} & \mathbb{E}[\mathbb{1}_{X_t \cap \Lambda_H \in A_H} \cdot \mathbb{1}_{X_{t+1} \cap \Lambda_B \in A_B} \cdot F_t \mid X_t \cap \Lambda_S; E] \\ &= \mathbb{E}[\mathbb{1}_{X_t \cap \Lambda_H \in A_H} \cdot \mathbb{E}[\mathbb{1}_{X_{t+1} \cap \Lambda_B \in A_B} \cdot F_t \mid X_t \cap (\Lambda_B \setminus \{\mathbf{v}\})^c; E] \mid X_t \cap \Lambda_S; E] \\ &= \mathbb{E}[\mathbb{1}_{X_t \cap \Lambda_H \in A_H} \cdot g \cdot h \mid X_t \cap \Lambda_S; E], \end{aligned} \quad (5)$$

where the first equality comes from the fact that $\sigma(X_t \cap \Lambda_S) \subseteq \sigma(X_t \cap (\Lambda_B \setminus \{\mathbf{v}\})^c)$. By Lemma 5.4, the map

$$(\omega, A_H) \mapsto \mu_{\lambda_{X_t(\omega) \cap \Lambda_S} \mathbb{1}_{(\Lambda_S)^c}}[\Lambda_H](A_H) \quad \text{for } \omega \in \Omega, A_H \in \mathfrak{A}_{\Lambda_H}$$

is a regular conditional distribution for $X_t \cap \Lambda_H$ given $X_t \cap \Lambda_S$ under the probability measure \mathbb{P}_E . To use this regular conditional probability to obtain an expression for the conditional expectation in (5), define functions $\hat{g}, \hat{h} : \mathcal{N}_{\Lambda_H} \times \mathcal{N}_{\Lambda_S} \rightarrow \mathbb{R}_{\geq 0}$ via

$$\begin{aligned} \hat{g}(\eta_1, \eta_2) &= \mu_{\lambda_{(\eta_1 \cup \eta_2) \cap (\Lambda_B)^c} \mathbb{1}_{\Lambda_B}}[\Lambda_B](A_B) \\ \hat{h}(\eta_1, \eta_2) &= C(S, \mathbf{v}, \eta_2) \cdot \frac{Z_{\Lambda_B}(\lambda_{(\eta_1 \cup \eta_2) \cap \Lambda_{\partial B}})}{Z_{\Lambda_B \setminus \{\mathbf{v}\}}(\lambda_{(\eta_1 \cup \eta_2) \cap \Lambda_{\partial B \cup \{\mathbf{v}\}})}}. \end{aligned}$$

Now, observe that for all $\omega \in \Omega$ it holds that $\hat{g}(X_t(\omega) \cap \Lambda_H, X_t(\omega) \cap \Lambda_S) = g(\omega)$ and $\hat{h}(X_t(\omega) \cap \Lambda_H, X_t(\omega) \cap \Lambda_S) = h(\omega)$. Thus, Theorem A.4 yields

$$\begin{aligned} & \mathbb{E}[\mathbb{1}_{X_t \cap \Lambda_H \in A_H} \cdot g \cdot h \mid X_t \cap \Lambda_S; E](\omega) \\ &= \mathbb{E}\left[\mathbb{1}_{X_t \cap \Lambda_H \in A_H} \cdot \hat{g}(X_t \cap \Lambda_H, X_t \cap \Lambda_S) \cdot \hat{h}(X_t \cap \Lambda_H, X_t \cap \Lambda_S) \mid X_t \cap \Lambda_S; E\right](\omega) \\ &= \int_{\mathcal{N}_{\Lambda_H}} \mathbb{1}_{\eta \in A_H} \cdot \hat{g}(\eta, X_t(\omega) \cap \Lambda_S) \cdot \hat{h}(\eta, X_t(\omega) \cap \Lambda_S) \mu_{\lambda_{X_t(\omega) \cap \Lambda_S} \mathbb{1}_{(\Lambda_S)^c}}[\Lambda_H](d\eta) \\ &= \frac{1}{Z_{(\Lambda_S)^c}(\lambda_{X_t(\omega) \cap \Lambda_S})} \sum_{n \geq 0} \frac{1}{n!} \int_{\Lambda_H^n} \mathbb{1}_{\eta_{\mathbf{x}} \in A_H} \cdot \hat{g}(\eta_{\mathbf{x}}, X_t(\omega) \cap \Lambda_S) \cdot \hat{h}(\eta_{\mathbf{x}}, X_t(\omega) \cap \Lambda_S) \\ &\quad \cdot (\lambda_{X_t(\omega) \cap \Lambda_S})^{\mathbf{x}} \cdot e^{-H(\mathbf{x})} \cdot Z_{(\Lambda_S)^c \cap (\Lambda_H)^c}(\lambda_{(X_t(\omega) \cap \Lambda_S) \cup \eta_{\mathbf{x}}}) d\mathbf{x} \\ &= \frac{1}{Z_{(\Lambda_S)^c}(\lambda_{X_t(\omega) \cap \Lambda_S})} \sum_{n \geq 0} \frac{1}{n!} \int_{\Lambda_H^n} \mathbb{1}_{\eta_{\mathbf{x}} \in A_H} \cdot \hat{g}(\eta_{\mathbf{x}}, X_t(\omega) \cap \Lambda_S) \cdot \hat{h}(\eta_{\mathbf{x}}, X_t(\omega) \cap \Lambda_S) \\ &\quad \cdot (\lambda_{X_t(\omega) \cap \Lambda_S})^{\mathbf{x}} \cdot e^{-H(\mathbf{x})} \cdot Z_{\Lambda_B \setminus \{\mathbf{v}\}}(\lambda_{(X_t(\omega) \cap \Lambda_S) \cup \eta_{\mathbf{x}}}) d\mathbf{x}. \end{aligned} \quad (6)$$

We proceed by simplifying (6). Observe that

$$\hat{h}(\eta_{\mathbf{x}}, X_t \cap \Lambda_S) = C(S, \mathbf{v}, X_t \cap \Lambda_S) \cdot \frac{Z_{\Lambda_B}(\lambda_{(\eta_{\mathbf{x}} \cup (X_t \cap \Lambda_S)) \cap \Lambda_{\partial B}})}{Z_{\Lambda_B \setminus \{\mathbf{v}\}}(\lambda_{(\eta_{\mathbf{x}} \cup (X_t \cap \Lambda_S)) \cap \Lambda_{\partial B \cup \{\mathbf{v}\}}})}$$

$$\begin{aligned}
&= C(S, \mathbf{v}, X_t \cap \Lambda_S) \cdot \frac{Z_{\Lambda_B} \left(\lambda_{(\eta_{\mathbf{x}} \cup (X_t \cap \Lambda_S)) \cap \Lambda_{(B)^c}} \right)}{Z_{\Lambda_{B \setminus \{\mathbf{v}\}}} \left(\lambda_{(\eta_{\mathbf{x}} \cup (X_t \cap \Lambda_S)) \cap (\Lambda_{B \setminus \{\mathbf{v}\}})^c} \right)} \\
&= C(S, \mathbf{v}, X_t \cap \Lambda_S) \cdot \frac{Z_{\Lambda_B} (\lambda_{\eta_{\mathbf{x}} \cup (X_t \cap \Lambda_R)})}{Z_{\Lambda_{B \setminus \{\mathbf{v}\}}} (\lambda_{\eta_{\mathbf{x}} \cup (X_t \cap \Lambda_S)})}.
\end{aligned}$$

Here, the second equality follows from Lemma C.2 and the fact that for all $\eta \in \mathcal{N}_\Lambda$ it holds that the distance between $(\eta \cap (\Lambda_B)^c) \Delta (\eta \cap \Lambda_{\partial B})$ and Λ_B , and the distance between $(\eta \cap (\Lambda_{B \setminus \{\mathbf{v}\}})^c) \Delta (\eta \cap \Lambda_{\partial B \cup \{\mathbf{v}\}})$ and $\Lambda_{B \setminus \{\mathbf{v}\}}$ are at least r . Moreover, the last equality follows from the fact that $S \cap (B)^c = S \setminus \{\mathbf{v}\} = R$, $\eta_{\mathbf{x}} \subseteq \Lambda_H \subseteq (\Lambda_B)^c \subseteq (\Lambda_{B \setminus \{\mathbf{v}\}})^c$ and $S \subseteq (B \setminus \{\mathbf{v}\})^c$. Further, note that

$$\begin{aligned}
\hat{g}(\eta_{\mathbf{x}}, X_t \cap \Lambda_S) &= \mu_{\lambda_{(\eta_{\mathbf{x}} \cup (X_t \cap \Lambda_S)) \cap (\Lambda_B)^c}} \mathbb{1}_{\Lambda_B} [\Lambda_B](A_B) \\
&= \mu_{\lambda_{\eta_{\mathbf{x}} \cup (X_t \cap \Lambda_R)}} \mathbb{1}_{\Lambda_B} [\Lambda_B](A_B) \\
&= \frac{\sum_{m \geq 0} \frac{1}{m!} \int_{\Lambda_B^m} \mathbb{1}_{\eta_{\mathbf{y}} \in A_B} \cdot (\lambda_{\eta_{\mathbf{x}} \cup (X_t \cap \Lambda_R)})^{\mathbf{y}} \cdot e^{-H(\mathbf{y})} d\mathbf{y}}{Z_{\Lambda_B} (\lambda_{\eta_{\mathbf{x}} \cup (X_t \cap \Lambda_R)})}.
\end{aligned}$$

Substituting both back into (6) and canceling $Z_{\Lambda_B} (\lambda_{\eta_{\mathbf{x}} \cup (X_t \cap \Lambda_R)})$ and $Z_{\Lambda_{B \setminus \{\mathbf{v}\}}} (\lambda_{\eta_{\mathbf{x}} \cup (X_t \cap \Lambda_S)})$ yields

$$\begin{aligned}
&\mathbb{E} \left[\mathbb{1}_{X_t \cap \Lambda_H \in A_H} \cdot \mathbb{1}_{X_{t+1} \cap \Lambda_B \in A_B} \cdot F_t \mid X_t \cap \Lambda_S; E \right] \\
&= \mathbb{E} \left[\mathbb{1}_{X_t \cap \Lambda_H \in A_H} \cdot g \cdot h \mid X_t \cap \Lambda_S; E \right] \\
&= \frac{C(S, \mathbf{v}, X_t \cap \Lambda_S)}{Z_{(\Lambda_S)^c} (\lambda_{X_t \cap \Lambda_S})} \sum_{n \geq 0} \frac{1}{n!} \int_{\Lambda_H^n} \mathbb{1}_{\eta_{\mathbf{x}} \in A_H} \cdot (\lambda_{X_t \cap \Lambda_S})^{\mathbf{x}} \cdot e^{-H(\mathbf{x})} \\
&\quad \cdot \left[\sum_{m \geq 0} \frac{1}{m!} \int_{\Lambda_B^m} \mathbb{1}_{\eta_{\mathbf{y}} \in A_B} \cdot (\lambda_{\eta_{\mathbf{x}} \cup (X_t \cap \Lambda_R)})^{\mathbf{y}} \cdot e^{-H(\mathbf{y})} d\mathbf{y} \right] d\mathbf{x}.
\end{aligned}$$

Moreover, note $\text{dist}(\Lambda_{\mathbf{v}}, \Lambda_H) \geq r$. Thus, it holds for $n \in \mathbb{N}_0$ and $\mathbf{x} = (x_1, \dots, x_n) \in \Lambda_H^n$ that

$$\begin{aligned}
(\lambda_{X_t \cap \Lambda_S})^{\mathbf{x}} &= \lambda^n e^{-\sum_{i=1}^n \sum_{z \in X_t \cap \Lambda_S} \phi(x_i, z)} \\
&= \lambda^n e^{-\sum_{i=1}^n \sum_{z \in X_t \cap \Lambda_R} \phi(x_i, z)} \\
&= (\lambda_{X_t \cap \Lambda_R})^{\mathbf{x}}.
\end{aligned}$$

Consequently, we have

$$\begin{aligned}
&\mathbb{E} \left[\mathbb{1}_{X_t \cap \Lambda_H \in A_H} \cdot \mathbb{1}_{X_{t+1} \cap \Lambda_B \in A_B} \cdot F_t \mid X_t \cap \Lambda_S; E \right] \\
&= \frac{C(S, \mathbf{v}, X_t \cap \Lambda_S)}{Z_{(\Lambda_S)^c} (\lambda_{X_t \cap \Lambda_S})} \sum_{n \geq 0} \frac{1}{n!} \int_{\Lambda_H^n} \mathbb{1}_{\eta_{\mathbf{x}} \in A_H} \cdot (\lambda_{X_t \cap \Lambda_R})^{\mathbf{x}} \cdot e^{-H(\mathbf{x})}
\end{aligned}$$

$$\cdot \left[\sum_{m \geq 0} \frac{1}{m!} \int_{\Lambda_B^m} \mathbb{1}_{\eta_{\mathbf{y}} \in A_B} \cdot (\lambda_{\eta_{\mathbf{x}} \cup (X_t \cap \Lambda_R)})^{\mathbf{y}} \cdot e^{-H(\mathbf{y})} d\mathbf{y} \right] d\mathbf{x}.$$

To shorten notation, let π_{Λ_B} denote the projection $\pi_{(\Lambda_R)^c, \Lambda_B}$ and write π_{Λ_H} for $\pi_{(\Lambda_R)^c, \Lambda_H}$. Since $H \cup B = (R)^c$, we obtain

$$\begin{aligned} & \sum_{n \geq 0} \frac{1}{n!} \int_{\Lambda_H^n} \mathbb{1}_{\eta_{\mathbf{x}} \in A_H} \cdot (\lambda_{X_t \cap \Lambda_R})^{\mathbf{x}} e^{-H(\mathbf{x})} \left[\sum_{m \geq 0} \frac{1}{m!} \int_{\Lambda_B^m} \mathbb{1}_{\eta_{\mathbf{y}} \in A_B} (\lambda_{\eta_{\mathbf{x}} \cup (X_t \cap \Lambda_R)})^{\mathbf{y}} e^{-H(\mathbf{y})} d\mathbf{y} \right] d\mathbf{x} \\ &= Z_{(\Lambda_R)^c}(\lambda_{X_t \cap \Lambda_R}) \cdot \mu_{\lambda_{X_t \cap \Lambda_R}} \mathbb{1}_{(\Lambda_R)^c} [(\Lambda_R)^c] \left(\pi_{\Lambda_H}^{-1}(A_H) \cap \pi_{\Lambda_B}^{-1}(A_B) \right). \end{aligned}$$

This finally yields the expression

$$\begin{aligned} & \mathbb{E}[\mathbb{1}_{X_t \cap \Lambda_H \in A_H} \cdot \mathbb{1}_{X_{t+1} \cap \Lambda_B \in A_B} \cdot F_t \mid X_t \cap \Lambda_S; E] \\ &= C(S, \mathbf{v}, X_t \cap \Lambda_S) \cdot \frac{Z_{(\Lambda_R)^c}(\lambda_{X_t \cap \Lambda_R})}{Z_{(\Lambda_S)^c}(\lambda_{X_t \cap \Lambda_S})} \\ & \quad \cdot \mu_{\lambda_{X_t \cap \Lambda_R}} \mathbb{1}_{(\Lambda_R)^c} [(\Lambda_R)^c] \left(\pi_{\Lambda_H}^{-1}(A_H) \cap \pi_{\Lambda_B}^{-1}(A_B) \right). \end{aligned}$$

Our next goal is to use Lemma A.1 to switch from \mathbb{P}_E to \mathbb{P}_{E_1} . To this end, note that

$$\begin{aligned} \mathbb{E}[F_t \mid X_t \cap \Lambda_S; E] &= \mathbb{E}[\mathbb{1}_{X_t \cap \Lambda_H \in \mathcal{N}_{\Lambda_H}} \cdot \mathbb{1}_{X_{t+1} \cap \Lambda_B \in \mathcal{N}_{\Lambda_B}} \cdot F_t \mid X_t \cap \Lambda_S; E] \\ &= C(S, \mathbf{v}, X_t \cap \Lambda_S) \cdot \frac{Z_{(\Lambda_R)^c}(\lambda_{X_t \cap \Lambda_R})}{Z_{(\Lambda_S)^c}(\lambda_{X_t \cap \Lambda_S})}. \end{aligned}$$

Thus, by Lemma A.1 it holds that

$$\begin{aligned} & \mathbb{E}[\mathbb{1}_{X_t \cap \Lambda_H \in A_H} \mathbb{1}_{X_{t+1} \cap \Lambda_B \in A_B} \mid X_t \cap \Lambda_S; E_1] \\ &= \frac{\mathbb{E}[\mathbb{1}_{X_t \cap \Lambda_H \in A_H} \cdot \mathbb{1}_{X_{t+1} \cap \Lambda_B \in A_B} \cdot F_t \mid X_t \cap \Lambda_S; E]}{\mathbb{E}[F_t \mid X_t \cap \Lambda_S; E]} \\ &= \mu_{\lambda_{X_t \cap \Lambda_R}} \mathbb{1}_{(\Lambda_R)^c} [(\Lambda_R)^c] \left(\pi_{\Lambda_H}^{-1}(A_H) \cap \pi_{\Lambda_B}^{-1}(A_B) \right). \end{aligned}$$

Moreover, given E_1 , it holds that $X_{t+1} \cap \Lambda_H = X_t \cap \Lambda_H$, which implies

$$\begin{aligned} & \mathbb{E}[\mathbb{1}_{X_{t+1} \cap \Lambda_H \in A_H} \cdot \mathbb{1}_{X_{t+1} \cap \Lambda_B \in A_B} \mid X_t \cap \Lambda_S; E_1] \\ &= \mu_{\lambda_{X_t \cap \Lambda_R}} \mathbb{1}_{(\Lambda_R)^c} [(\Lambda_R)^c] \left(\pi_{\Lambda_H}^{-1}(A_H) \cap \pi_{\Lambda_B}^{-1}(A_B) \right) \end{aligned} \quad (7)$$

We now generalize (7) to arbitrary events $A \in \mathfrak{A}_{(\Lambda_R)^c}$. To this end, note that $\Lambda_H \cup \Lambda_B = (\Lambda_R)^c$. Thus, Lemma C.1 shows that events of the form

$$\left\{ \pi_{\Lambda_H}^{-1}(A_H) \cap \pi_{\Lambda_B}^{-1}(A_B) \mid A_H \in \mathfrak{A}_{\Lambda_H}, A_B \in \mathfrak{A}_{\Lambda_B} \right\}$$

are a generating π -system for $\mathfrak{R}_{(\Lambda_R)^c}$. Since further $\mu_{\lambda_{X_t \cap \Lambda_R} \mathbf{1}_{(\Lambda_R)^c}}[(\Lambda_R)^c]$ is a probability distribution on $(\mathcal{N}_{(\Lambda_R)^c}, \mathfrak{R}_{(\Lambda_R)^c})$, Lemma A.5 implies that for all $A \in \mathfrak{R}_{\Lambda_R}$

$$\mathbb{E}[\mathbf{1}_{X_{t+1} \cap (\Lambda_R)^c \in A} \mid X_t \cap \Lambda_S; E_1] = \mu_{\lambda_{X_t \cap \Lambda_R} \mathbf{1}_{(\Lambda_R)^c}}[(\Lambda_R)^c](A). \quad (8)$$

Note that the right-hand side of (8) is $\sigma(X_t \cap \Lambda_R)$ -measurable. As $\sigma(X_t \cap \Lambda_R) \subset \sigma(X_t \cap \Lambda_S)$, this implies that it is also a version of the conditional expectation $\mathbb{E}[\mathbf{1}_{X_{t+1} \cap (\Lambda_R)^c \in A} \mid X_t \cap \Lambda_R; E_1]$. Finally, observing that, given E_1 , it holds that $X_{t+1} \cap \Lambda_R = X_t \cap \Lambda_R$ we obtain

$$\mathbb{E}[\mathbf{1}_{X_{t+1} \cap (\Lambda_R)^c \in A} \mid X_{t+1} \cap \Lambda_R; E_1] = \mu_{\lambda_{X_{t+1} \cap \Lambda_R} \mathbf{1}_{(\Lambda_R)^c}}[(\Lambda_R)^c](A). \quad \square$$

The next lemma is the counterpart of Lemma 5.5 for the case $F_t = 0$.

Lemma 5.6. *Under the assumptions of Lemma 5.4, let $E_0 = \{\mathcal{U}_t = S, \mathbf{u}_t = \mathbf{v}, F_t = 0\}$ for some $S \in 2^{\mathcal{V}} \setminus \{\emptyset\}$ and $\mathbf{v} \in S$ such that $\mathbb{P}[E_0] > 0$, and set $B = B(S, \mathbf{v}, \ell)$ and $R = S \cup B$. For all $A \in \mathfrak{R}_{(\Lambda_R)^c}$ it holds that*

$$\mathbb{E}[\mathbf{1}_{X_{t+1} \cap (\Lambda_R)^c \in A} \mid X_{t+1} \cap \Lambda_R; E_0] = \mu_{\lambda_{X_{t+1} \cap \Lambda_R} \mathbf{1}_{(\Lambda_R)^c}}[(\Lambda_R)^c](A).$$

Proof. Let $E = \{\mathcal{U}_t = S, \mathbf{u}_t = \mathbf{v}\}$ for S and \mathbf{v} as in the definition of E_0 . Our first step is to show that for all $A \in \mathfrak{R}_{(\Lambda_R)^c}$ it holds that

$$\mathbb{E}[\mathbf{1}_{X_t \cap (\Lambda_R)^c \in A} \mid X_t \cap \Lambda_R; E] = \mu_{\lambda_{X_t \cap \Lambda_R} \mathbf{1}_{(\Lambda_R)^c}}[(\Lambda_R)^c](A). \quad (9)$$

Define $f : \Omega \rightarrow [0, 1]$ to be the right-hand side of the equality above. Since f is $\sigma(X_t \cap \Lambda_R)$ -measurable, it suffices to prove that

$$\mathbb{E}[\mathbf{1}_{X_t \cap \Lambda_R \in D} \cdot f \mid E] = \mathbb{E}[\mathbf{1}_{X_t \cap \Lambda_R \in D} \cdot \mathbf{1}_{X_t \cap (\Lambda_R)^c \in A} \mid E] \quad (10)$$

for all events $D \in \mathfrak{R}_{\Lambda_R}$. For this, it suffices to prove this statement for some π -system of events $\mathcal{D} \subseteq \mathfrak{R}_{\Lambda_R}$ that generates \mathfrak{R}_{Λ_R} . Write $\pi_{\Lambda_{\partial B \setminus S}}$ for the projection $\pi_{\Lambda_R, \Lambda_{\partial B \setminus S}}$, and write π_{Λ_S} for $\pi_{\Lambda_R, \Lambda_S}$. Since R is the union of the disjoint sets S and $\partial B \setminus S$, Lemma C.1 shows that such a generating π -system for \mathfrak{R}_{Λ_R} is given by

$$\mathcal{D} = \left\{ \pi_{\Lambda_{\partial B \setminus S}}^{-1}(D_1) \cap \pi_{\Lambda_S}^{-1}(D_2) \mid D_1 \in \mathfrak{R}_{\Lambda_{\partial B \setminus S}}, D_2 \in \mathfrak{R}_{\Lambda_S} \right\}.$$

Let $D_1 \in \mathfrak{R}_{\Lambda_{\partial B \setminus S}}$ and $D_2 \in \mathfrak{R}_{\Lambda_S}$. Writing

$$\mathbb{E}[\mathbf{1}_{X_t \cap \Lambda_{\partial B \setminus S} \in D_1} \cdot \mathbf{1}_{X_t \cap S \in D_2} \cdot f \mid E] = \mathbb{E}[\mathbf{1}_{X_t \cap S \in D_2} \cdot \mathbb{E}[\mathbf{1}_{X_t \cap \Lambda_{\partial B \setminus S} \in D_1} \cdot f \mid X_t \cap \Lambda_S; E] \mid E]$$

shows that, for proving (10) for all $D \in \mathcal{D}$ (and by extension for all $D \in \mathfrak{R}_{\Lambda_R}$), it suffices to prove

$$\mathbb{E}[\mathbf{1}_{X_t \cap \Lambda_{\partial B \setminus S} \in D_1} \cdot f \mid X_t \cap \Lambda_S; E] = \mathbb{E}[\mathbf{1}_{X_t \cap \Lambda_{\partial B \setminus S} \in D_1} \cdot \mathbf{1}_{X_t \cap (\Lambda_R)^c \in A} \mid X_t \cap \Lambda_S; E].$$

To this end, define $\hat{f} : \mathcal{N}_{\Lambda_{\partial B \setminus S}} \times \mathcal{N}_{\Lambda_S} \rightarrow [0, 1]$ by

$$\hat{f}(\eta_1, \eta_2) = \mu_{\lambda_{\eta_1 \cup \eta_2}} \mathbb{1}_{(\Lambda_R)^c} [(\Lambda_R)^c](A)$$

and note that $\hat{f}(X_t(\omega) \cap \Lambda_{\partial B \setminus S}, X_t(\omega) \cap \Lambda_S) = f(\omega)$ for all $\omega \in \Omega$. Therefore, we have

$$\begin{aligned} & \mathbb{E} \left[\mathbb{1}_{X_t \cap \Lambda_{\partial B \setminus S} \in D_1} \cdot f \mid X_t \cap \Lambda_S; E \right] \\ &= \mathbb{E} \left[\mathbb{1}_{X_t \cap \Lambda_{\partial B \setminus S} \in D_1} \cdot \hat{f}(X_t \cap \Lambda_{\partial B \setminus S}, X_t \cap \Lambda_S) \mid X_t \cap \Lambda_S; E \right]. \end{aligned}$$

We proceed by deriving an explicit expression for the right-hand side. By Lemma 5.4 and the assumption of the lemma, we know that

$$(\omega, \cdot) \mapsto \mu_{\lambda_{X_t(\omega) \cap \Lambda_S}} \mathbb{1}_{(\Lambda_S)^c} [\Lambda_{\partial B \setminus S}](\cdot)$$

is regular conditional distribution for $X_t \cap \Lambda_{\partial B \setminus S}$ given $X_t \cap \Lambda_S$ under \mathbb{P}_E . Thus, using Theorem A.4, we obtain

$$\begin{aligned} & \mathbb{E} \left[\mathbb{1}_{X_t \cap \Lambda_{\partial B \setminus S} \in D_1} \cdot \hat{f}(X_t \cap \Lambda_{\partial B \setminus S}, X_t \cap \Lambda_S) \mid X_t \cap \Lambda_S; E \right](\omega) \\ &= \int_{D_1} \hat{f}(\eta, X_t(\omega) \cap \Lambda_S) \mu_{\lambda_{X_t(\omega) \cap \Lambda_S}} \mathbb{1}_{(\Lambda_S)^c} [\Lambda_{\partial B \setminus S}](d\eta) \\ &= \frac{1}{Z_{(\Lambda_S)^c}(\lambda_{X_t(\omega) \cap \Lambda_S})} \sum_{n \geq 0} \frac{1}{n!} \int_{\Lambda_{\partial B \setminus S}^n} \mathbb{1}_{\eta_{\mathbf{x}} \in D_1} \cdot \hat{f}(\eta_{\mathbf{x}}, X_t(\omega) \cap \Lambda_S) \\ & \quad \cdot (\lambda_{X_t(\omega) \cap \Lambda_S})^{\mathbf{x}} \cdot e^{-H(\mathbf{x})} \cdot Z_{(\Lambda_{\partial B \setminus S})^c \cap (\Lambda_S)^c}(\lambda_{(X_t(\omega) \cap \Lambda_S) \cup \eta_{\mathbf{x}}}) d\mathbf{x} \\ &= \frac{1}{Z_{(\Lambda_S)^c}(\lambda_{X_t(\omega) \cap \Lambda_S})} \sum_{n \geq 0} \frac{1}{n!} \int_{\Lambda_{\partial B \setminus S}^n} \mathbb{1}_{\eta_{\mathbf{x}} \in D_1} \cdot \hat{f}(\eta_{\mathbf{x}}, X_t(\omega) \cap \Lambda_S) \\ & \quad \cdot (\lambda_{X_t(\omega) \cap \Lambda_S})^{\mathbf{x}} \cdot e^{-H(\mathbf{x})} \cdot Z_{(\Lambda_R)^c}(\lambda_{(X_t(\omega) \cap \Lambda_S) \cup \eta_{\mathbf{x}}}) d\mathbf{x} \end{aligned}$$

for \mathbb{P}_E -almost all $\omega \in \Omega$. Further, note that

$$\begin{aligned} & \hat{f}(\eta_{\mathbf{x}}, X_t(\omega) \cap \Lambda_S) \\ &= \mu_{\lambda_{(X_t(\omega) \cap \Lambda_S) \cup \eta_{\mathbf{x}}}} \mathbb{1}_{(\Lambda_R)^c} [(\Lambda_R)^c](A) \\ &= \frac{1}{Z_{(\Lambda_R)^c}(\lambda_{(X_t(\omega) \cap \Lambda_S) \cup \eta_{\mathbf{x}}})} \sum_{m \geq 0} \frac{1}{m!} \int_{((\Lambda_R)^c)^m} \mathbb{1}_{\eta_{\mathbf{y}} \in A} \cdot (\lambda_{(X_t(\omega) \cap \Lambda_S) \cup \eta_{\mathbf{x}}})^{\mathbf{y}} \cdot e^{-H(\mathbf{y})} d\mathbf{y}. \end{aligned}$$

Thus, after canceling $Z_{(\Lambda_R)^c}(\lambda_{(X_t(\omega) \cap \Lambda_S) \cup \eta_{\mathbf{x}}})$, we obtain

$$\begin{aligned} & \mathbb{E} \left[\mathbb{1}_{X_t \cap \Lambda_{\partial B \setminus S} \in D_1} \cdot \hat{f}(X_t \cap \Lambda_{\partial B \setminus S}, X_t \cap \Lambda_S) \mid X_t \cap \Lambda_S; E \right](\omega) \\ &= \frac{1}{Z_{(\Lambda_S)^c}(\lambda_{X_t(\omega) \cap \Lambda_S})} \sum_{n \geq 0} \frac{1}{n!} \int_{\Lambda_{\partial B \setminus S}^n} \mathbb{1}_{\eta_{\mathbf{x}} \in D_1} \cdot (\lambda_{X_t(\omega) \cap \Lambda_S})^{\mathbf{x}} \cdot e^{-H(\mathbf{x})} \end{aligned}$$

$$\begin{aligned}
& \cdot \left[\sum_{m \geq 0} \frac{1}{m!} \int_{((\Lambda_R)^c)^m} \mathbb{1}_{\eta_{\mathbf{y}} \in A} \cdot (\lambda_{(X_t(\omega) \cap \Lambda_S) \cup \eta_{\mathbf{x}}})^{\mathbf{y}} \cdot e^{-H(\mathbf{y})} d\mathbf{y} \right] d\mathbf{x} \\
& = \mu_{\lambda_{X_t(\omega) \cap \Lambda_S} \mathbb{1}_{(\Lambda_S)^c}} [(\Lambda_S)^c] \left(\pi_{\Lambda_{\partial B \setminus S}}^{-1}(D_1) \cap \pi_{(\Lambda_R)^c}^{-1}(A) \right) \\
& = \mathbb{E} \left[\mathbb{1}_{X_t \cap \Lambda_{\partial B \setminus S} \in D_1} \cdot \mathbb{1}_{X_t \cap (\Lambda_R)^c \in A} \mid X_t \cap \Lambda_S; E \right] (\omega)
\end{aligned}$$

for \mathbb{P}_E -almost all $\omega \in \Omega$, where the last equality is once again due to Lemma 5.4. This shows (10) and consequently (9).

In the next step, we aim for changing the probability measure from \mathbb{P}_E to \mathbb{P}_{E_0} . To achieve this, note that by definition F_t does only depend on X_t via $X_t \cap \Lambda_R$. Thus, it holds that

$$\mathbb{E}[\mathbb{1}_{F_t=0} \mid X_t \cap \Lambda_R; E] = \mathbb{E}[\mathbb{1}_{F_t=0} \mid X_t \cap \Lambda_R, X_t \cap (\Lambda_R)^c; E].$$

Therefore, Lemma A.3 implies

$$\begin{aligned}
\mathbb{E}[\mathbb{1}_{X_t \cap (\Lambda_R)^c \in A} \mid X_t \cap \Lambda_R; E_0] &= \mathbb{E}[\mathbb{1}_{X_t \cap (\Lambda_R)^c \in A} \mid X_t \cap \Lambda_R; E] \\
&= \mu_{\lambda_{X_t \cap \Lambda_R} \mathbb{1}_{(\Lambda_R)^c}} [(\Lambda_R)^c](A).
\end{aligned}$$

Finally, observe that, given E_0 , it holds that $X_{t+1} = X_t$. Thus, we have

$$\mathbb{E}[\mathbb{1}_{X_{t+1} \cap (\Lambda_R)^c \in A} \mid X_{t+1} \cap \Lambda_R; E_0] = \mu_{\lambda_{X_{t+1} \cap \Lambda_R} \mathbb{1}_{(\Lambda_R)^c}} [(\Lambda_R)^c](A)$$

as desired. \square

Combining Lemmas 5.5 and 5.6, we now prove Lemma 5.3.

Proof of Lemma 5.3. First, note that it suffices to show that, for all $t \in \mathbb{N}_0$ and $S \subseteq \mathcal{V}$ with $\mathbb{P}[\mathcal{U}_t = S] > 0$, it holds that $\mu_{\lambda_{X_t \cap \Lambda_S} \mathbb{1}_{(\Lambda_S)^c}} [(\Lambda_S)^c](A)$ is a version of the conditional expectation $\mathbb{E}[\mathbb{1}_{X_t \cap (\Lambda_S)^c \in A} \mid X_t \cap \Lambda_S; \mathcal{U}_t = S]$. The second part of the statement then follows as $\mu_{\lambda_{\eta} \mathbb{1}_{(\Lambda_S)^c}} [(\Lambda_S)^c]$ is a probability distribution on $(\mathcal{N}_{(\Lambda_S)^c}, \mathfrak{R}_{(\Lambda_S)^c})$ for all $\eta \in \mathcal{N}_{\Lambda_S}$.

We proceed by proving our claim by an induction over t . For $t = 0$, the statement is trivially true since $\mathcal{U}_0 = \mathcal{V}$ and $(\Lambda_{\mathcal{V}})^c = \emptyset$.

Now, assume the lemma holds for some fixed $t \in \mathbb{N}$ and let $R \subseteq \mathcal{V}$ be such that $\mathbb{P}[\mathcal{U}_{t+1} = R] > 0$. We start with the case $R \neq \emptyset$. We consider two sets of event \mathcal{C}_0 and \mathcal{C}_1 , where \mathcal{C}_i for $i \in \{0, 1\}$ consists of all events of the form $C \{ \mathcal{U}_t = S, \mathbf{u}_t = \mathbf{v}, F_t = i \}$ with $S \in 2^{\mathcal{V}} \setminus \{\emptyset\}$ and $\mathbf{v} \in S$ such that $\mathbb{P}[C \cap \{ \mathcal{U}_{t+1} = R \}] > 0$. Set $\mathcal{C} = \mathcal{C}_0 \cup \mathcal{C}_1$ and note that all events in \mathcal{C} are pairwise disjoint. Moreover, it is easy to check that

$$\mathbb{P}_{\{ \mathcal{U}_{t+1} = R \}} \left[\bigcup_{C \in \mathcal{C}} C \right] = \mathbb{P} \left[\bigcup_{C \in \mathcal{C}} C \mid \mathcal{U}_{t+1} = R \right] = 1.$$

Thus, given we show that

$$\mathbb{E}[\mathbb{1}_{X_{t+1} \cap (\Lambda_R)^c \in A} \mid X_t \cap \Lambda_R; \mathcal{C}] = \mu_{\lambda_{X_{t+1} \cap \Lambda_R} \mathbb{1}_{(\Lambda_R)^c}} [(\Lambda_R)^c](A) \quad (11)$$

for all $C \in \mathcal{C}$, Lemma A.2 implies that $\mu_{\lambda_{X_{t+1} \cap \Lambda_R} \mathbb{1}_{(\Lambda_R)^c}}[(\Lambda_R)^c](A)$ is also a version of $\mathbb{E}[\mathbb{1}_{X_{t+1} \cap (\Lambda_R)^c \in A} \mid X_t \cap \Lambda_R; \mathcal{U}_{t+1} = R]$ as desired. Suppose that $C \in \mathcal{C}_1$. Then C must have the form $\{\mathcal{U}_t = S, \mathbf{u}_t = \mathbf{v}, F_t = 1\}$ for some S and \mathbf{v} with $R = S \setminus \{\mathbf{v}\}$. Thus, using the induction hypothesis and applying Lemma 5.5 proves (11). Otherwise, if $C \in \mathcal{C}_0$, then C is of the form $\{\mathcal{U}_t = S, \mathbf{u}_t = \mathbf{v}, F_t = 0\}$ with $R = S \cup \partial B(S, \mathbf{v}, \ell)$. Using the induction hypothesis and Lemma 5.6 shows (11).

It remains to consider the case $R = \emptyset$. We construct \mathcal{C}_0 and \mathcal{C}_1 as before, but we set $\mathcal{C} = \mathcal{C}_0 \cup \mathcal{C}_1 \cup \{\{\mathcal{U}_t = \emptyset\}\}$. Again, by Lemma A.2 it suffices to argue (11) for all $C \in \mathcal{C}$. The cases $C \in \mathcal{C}_1$ and $C \in \mathcal{C}_0$ are handled as before and we focus on $C = \{\mathcal{U}_t = \emptyset\}$. By our definition of the process, we fixed X_t and \mathcal{U}_t to remain constant once $\mathcal{U}_t = \emptyset$. Thus, (11) follows directly from the induction hypothesis, concluding the proof. \square

6 Strong spatial mixing and success probabilities of Bayes filters

Recall that for Algorithm 1 to terminate rapidly, we need to ensure that the success probability of the Bayes filter is close to 1. In this section, we prove a general statement that allows us to control the success probabilities of the Bayes filters we will construct in the upcoming sections. Readers only interested in the actual construction of the Bayes filter may skip this section for now and return to it later for the running time analysis.

The main technical lemma of this section states that, under strong spatial mixing, a certain fraction of partition functions that is central for the construction of our Bayes filters can be brought arbitrarily close to 1 by increasing the update radius ℓ .

Lemma 6.1. *Let $S \subseteq \mathcal{V}$ be non-empty, $\mathbf{v} \in S$, $\ell \in \mathbb{N}$, and set $B = B(S, \mathbf{v}, \ell)$ and $H = (S \cup B)^c$. Suppose ϕ exhibits (a, b) -strong spatial mixing up to λ . Then, for all feasible $\eta \in \mathcal{N}_\Lambda$ and all $\xi_1, \xi_2 \in \mathcal{N}_{\Lambda_H}$ it holds that*

$$\begin{aligned} \exp\left(-a3^d r^d e^{2br} \left(\lambda r^d + e^{\lambda 3^d r^d}\right) e^{-br\ell}\right) &\leq \frac{Z_{\Lambda_{B \setminus \{\mathbf{v}\}}}(\lambda_{\xi_1 \cup (\eta \cap \Lambda_S)})}{Z_{\Lambda_B}(\lambda_{\xi_1 \cup (\eta \cap \Lambda_{S \setminus \{\mathbf{v}\}})})} \cdot \frac{Z_{\Lambda_B}(\lambda_{\xi_2 \cup (\eta \cap \Lambda_{S \setminus \{\mathbf{v}\}})})}{Z_{\Lambda_{B \setminus \{\mathbf{v}\}}}(\lambda_{\xi_2 \cup (\eta \cap \Lambda_S)})} \\ &\leq \exp\left(a3^d r^d e^{2br} \left(\lambda r^d + e^{\lambda 3^d r^d}\right) e^{-br\ell}\right). \end{aligned}$$

Before we prove Lemma 6.1, we briefly sketch how it helps control the success probability of the Bayes filter. Recall Definition 4.1 and assume we would directly use

$$C(S, \mathbf{v}, \eta) = \inf_{\substack{\xi \in \mathcal{N}_{\Lambda_H} \\ \xi \cup (\eta \cap \Lambda_S) \text{ is feasible}}} \left\{ \frac{Z_{\Lambda_{B \setminus \{\mathbf{v}\}}}(\lambda_{\xi \cup (\eta \cap \Lambda_S)})}{Z_{\Lambda_B}(\lambda_{\xi \cup (\eta \cap \Lambda_{S \setminus \{\mathbf{v}\}})})} \right\} \quad (12)$$

as Bayes filter correction. Assuming $X_t = \eta$, $\mathcal{U}_t = S$ and $\mathbf{u}_t = \mathbf{v} \in S$, Lemma C.2 yields that the probability that $F_t = 1$ is

$$\begin{aligned} & \inf_{\substack{\xi \in \mathcal{N}_{\Lambda_H} \\ \xi \cup (\eta \cap \Lambda_S) \text{ is feasible}}} \left\{ \frac{Z_{\Lambda_{B \setminus \{\mathbf{v}\}}}(\lambda_{\xi \cup (\eta \cap \Lambda_S)})}{Z_{\Lambda_B}(\lambda_{\xi \cup (\eta \cap \Lambda_S \setminus \{\mathbf{v}\})})} \right\} \cdot \frac{Z_{\Lambda_B}(\lambda_{\eta \cap \Lambda_{\partial B}})}{Z_{\Lambda_{B \setminus \{\mathbf{v}\}}}(\lambda_{\eta \cap \Lambda_{\partial B \cup \{\mathbf{v}\}}})} \\ &= \inf_{\substack{\xi \in \mathcal{N}_{\Lambda_H} \\ \xi \cup (\eta \cap \Lambda_S) \text{ is feasible}}} \left\{ \frac{Z_{\Lambda_{B \setminus \{\mathbf{v}\}}}(\lambda_{\xi \cup (\eta \cap \Lambda_S)})}{Z_{\Lambda_B}(\lambda_{\xi \cup (\eta \cap \Lambda_S \setminus \{\mathbf{v}\})})} \cdot \frac{Z_{\Lambda_B}(\lambda_{(\eta \cap \Lambda_H) \cup (\eta \cap \Lambda_S \setminus \{\mathbf{v}\})})}{Z_{\Lambda_{B \setminus \{\mathbf{v}\}}}(\lambda_{(\eta \cap \Lambda_H) \cup (\eta \cap \Lambda_S)})} \right\}. \end{aligned}$$

Applying Lemma 6.1 with $\xi_1 = \xi$ and $\xi_2 = \eta \cap \Lambda_H$ allows us to lower bound the probability that $F_t = 1$ by $\exp\left(-a3^d r^d e^{2br} \left(\lambda r^d + e^{\lambda 3^d r^d}\right) e^{-br\ell}\right)$. Thus, by increasing the update radius ℓ we could bring the success probability of the Bayes filter arbitrary close to 1. While we will not use exactly (12) as Bayes filter correction, we can apply Lemma 6.1 in a similar fashion when using a suitable approximation. More on that in Section 7 and Section 8.

To prove Lemma 6.1, we first show that strong spatial mixing implies correlation decay in terms of k -point density functions. The converse of this statement was previously shown in [58]. We then use an identity from [58, 56] to derive our lemma.

6.1 Strong spatial mixing and point density functions

For $k \in \mathbb{N}$ and $\mathbf{x} \in \Lambda^k$ the k -point density function of a Gibbs point process μ_λ is defined as

$$\rho_\lambda(\mathbf{x}) = \lambda^{\mathbf{x}} \frac{Z_\Lambda(\lambda_{\mathbf{x}})}{Z_\Lambda(\lambda)} e^{-H(\mathbf{x})}.$$

Recall Definition 1.1 and note that, for a measurable space (Ω, \mathcal{A}) and probability measures P and Q , an equivalent definition of total variation distance is

$$d_{TV}(P, Q) = \sup_{f: \Omega \rightarrow [-1, 1]} |\mathbb{E}_P[f] - \mathbb{E}_Q[f]|,$$

where the supremum is taken over measurable functions. Using this definition we obtain the following statement.

Lemma 6.2. *Let ϕ be a repulsive pair potential of range r and let $\lambda, a, b \in \mathbb{R}_{>0}$ such that ϕ exhibits (a, b) -strong spatial mixing up to λ . Let $\Lambda \subset \mathbb{R}^d$ be bounded and measurable, and consider activity functions $\lambda, \lambda' < \lambda$. For all measurable $\Lambda' \subseteq \Lambda$ disjoint from $\text{supp}(\lambda - \lambda')$, all $k \in \mathbb{N}$ and all $\mathbf{x} \in \Lambda^k$ it holds that*

$$|\rho_\lambda(\mathbf{x}) - \rho_{\lambda'}(\mathbf{x})| \leq \lambda^{\mathbf{x}} e^{-H(\mathbf{x})} a \left| \Lambda'^{(r)} \right| e^{br} e^{-b \cdot \text{dist}(\Lambda', \text{supp}(\lambda - \lambda'))},$$

where $\Lambda'^{(r)} = \{y \in \Lambda \mid \text{dist}(y, \Lambda') < r\}$, and ρ_λ and $\rho_{\lambda'}$ are the k -point densities on Λ for the potential ϕ and activity functions λ and λ' .

Proof. For every $k \in \mathbb{N}$ and $\mathbf{x} = (x_1, \dots, x_k) \in \Lambda'^k$ define $f_{\mathbf{x}} : \mathcal{N}_{\Lambda} \rightarrow [0, 1]$ by

$$f_{\mathbf{x}}(\eta) = e^{-\sum_{i \in [k]} \sum_{y \in \eta} \phi(x_i, y)}.$$

By definition, it holds that

$$\begin{aligned} \rho_{\lambda}(\mathbf{x}) &= \lambda^{\mathbf{x}} \frac{Z_{\Lambda}(\lambda_{\mathbf{x}})}{Z_{\Lambda}(\lambda)} e^{-H(\mathbf{x})} \\ &= \lambda^{\mathbf{x}} e^{-H(\mathbf{x})} \mathbb{E}_{\mu_{\lambda}}[f_{\mathbf{x}}]. \end{aligned}$$

Since the range of ϕ is bounded by r , it holds that $f_{\mathbf{x}}$ is local on $\Lambda'^{(r)} = \{y \in \Lambda \mid \text{dist}(y, \Lambda') < r\}$. That is, for every $\eta \in \mathcal{N}_{\Lambda}$ it holds that $f_{\mathbf{x}}(\eta) = f_{\mathbf{x}}(\eta \cap \Lambda'^{(r)})$. Consequently, we have $\mathbb{E}_{\mu_{\lambda}}[f_{\mathbf{x}}] = \mathbb{E}_{\mu_{\lambda}[\Lambda'^{(r)}]}[f_{\mathbf{x}}]$. Applying the same reasoning to $\rho_{\lambda'}(\mathbf{x})$ and observing that $\lambda = \lambda'$ on Λ' yields

$$\begin{aligned} |\rho_{\lambda}(\mathbf{x}) - \rho_{\lambda'}(\mathbf{x})| &= \lambda^{\mathbf{x}} e^{-H(\mathbf{x})} \left| \mathbb{E}_{\mu_{\lambda}[\Lambda'^{(r)}]}[f_{\mathbf{x}}] - \mathbb{E}_{\mu_{\lambda'}[\Lambda'^{(r)}]}[f_{\mathbf{x}}] \right| \\ &\leq \lambda^{\mathbf{x}} e^{-H(\mathbf{x})} d_{TV} \left(\mu_{\lambda}[\Lambda'^{(r)}], \mu_{\lambda'}[\Lambda'^{(r)}] \right), \end{aligned}$$

where the inequality follows from the definition of the total variation distance given above and the fact that $f_{\mathbf{x}}$ has domain $[0, 1]$. Finally, applying (a, b) -strong spatial mixing and noting that $\text{dist}(\Lambda'^{(r)}, \text{supp}(\lambda - \lambda')) \geq \text{dist}(\Lambda', \text{supp}(\lambda - \lambda')) - r$ concludes the proof. \square

Remark 6.3. Note that, without fixing a particular region $\Lambda' \subseteq \Lambda$ that contains x_1, \dots, x_k in Lemma 6.2, we can always set $\Lambda' = \{x_1, \dots, x_k\}$, which yields $|\Lambda'^{(r)}| \leq kv_d r^d$ with v_d being the volume of a unit ball in d dimensions.

The following multiplicative bound for k -point densities with different activity functions follows immediately.

Corollary 6.4. Consider the setting of Lemma 6.2. It holds that

$$\rho_{\lambda}(\mathbf{x}) \leq \left(1 + a |\Lambda'^{(r)}| e^{br + \lambda |\Lambda'^{(r)}|} e^{-b \cdot \text{dist}(\Lambda', \text{supp}(\lambda - \lambda'))} \right) \rho_{\lambda'}(\mathbf{x}).$$

Proof. Since $\lambda = \lambda'$ on Λ' , it holds that $\lambda^{\mathbf{x}} e^{-H(\mathbf{x})} = \lambda'^{\mathbf{x}} e^{-H(\mathbf{x})}$. If $\lambda^{\mathbf{x}} e^{-H(\mathbf{x})} = 0$, then the desired inequality holds trivially since both sides are 0.

Assume $\lambda^{\mathbf{x}} e^{-H(\mathbf{x})} > 0$. Defining $f_{\mathbf{x}}$ as in the proof of Lemma 6.2 and following the same arguments we have

$$\rho_{\lambda'}(\mathbf{x}) = \lambda'^{\mathbf{x}} e^{-H(\mathbf{x})} \mathbb{E}_{\mu_{\lambda'}[\Lambda'^{(r)}]}[f_{\mathbf{x}}] = \lambda^{\mathbf{x}} e^{-H(\mathbf{x})} \mathbb{E}_{\mu_{\lambda'}[\Lambda'^{(r)}]}[f_{\mathbf{x}}].$$

Observe that $f_{\mathbf{x}}$ is non-negative and $f_{\mathbf{x}}(\emptyset) = 1$. Combined with Poisson domination, we obtain

$$\mathbb{E}_{\mu_{\lambda'}[\Lambda'^{(r)}]}[f_{\mathbf{x}}] \geq \mu_{\lambda'}[\Lambda'^{(r)}](\{\emptyset\}) \geq e^{-\lambda |\Lambda'^{(r)}|},$$

which implies $\rho_{\lambda'}(\mathbf{x}) \geq \lambda^{\mathbf{x}} e^{-H(\mathbf{x})} e^{-\lambda^{|\Lambda'(r)|}}$. In particular, we have $\rho_{\lambda'}(\mathbf{x}) > 0$, and applying Lemma 6.2 yields

$$\begin{aligned} \rho_{\lambda}(\mathbf{x}) &\leq \left(1 + \frac{|\rho_{\lambda}(\mathbf{x}) - \rho_{\lambda'}(\mathbf{x})|}{\rho_{\lambda'}(\mathbf{x})}\right) \rho_{\lambda'}(\mathbf{x}) \\ &\leq \left(1 + a \left|\Lambda'(r)\right| e^{br+\lambda^{|\Lambda'(r)|}} e^{-b \cdot \text{dist}(\Lambda', \text{supp}(\lambda - \lambda'))}\right) \rho_{\lambda'}(\mathbf{x}). \end{aligned} \quad \square$$

6.2 Proof of Lemma 6.1

To prove the main lemma of the section, we use the following identity by Michelen and Perkins [58, 56].

Lemma 6.5 ([58, Lemma 12]). *Let $\Lambda' \subseteq \Lambda$ be measurable. Fix a point in $z \in \Lambda'$ and, for any given activity function λ and any point $y \in \Lambda$, let*

$$\widehat{\lambda}_y(w) = \begin{cases} 0 & \text{if } \widetilde{\text{dist}}(z, w) < \widetilde{\text{dist}}(z, y) \\ \lambda(w) & \text{otherwise} \end{cases},$$

where $\widetilde{\text{dist}}(u, v) = \text{dist}(u, v) + \text{diam}(\Lambda') \cdot \mathbb{1}_{\{u, v\} \not\subseteq \Lambda'}$ for $u, v \in \mathbb{R}^d$. For all $k \in \mathbb{N}$ and $\mathbf{x} \in \Lambda'^k$, it holds that

$$\lambda^{\mathbf{x}} \frac{Z_{(\Lambda')^c}(\lambda_{\mathbf{x}})}{Z_{\Lambda}(\lambda)} e^{-H(\mathbf{x})} = \rho_{\lambda}(\mathbf{x}) \exp\left(-\int_{\Lambda'} \rho_{(\widehat{\lambda}_{\mathbf{x}})_y}(y) dy\right).$$

Moreover, we have

$$\frac{Z_{(\Lambda')^c}(\lambda)}{Z_{\Lambda}(\lambda)} = \exp\left(-\int_{\Lambda'} \rho_{\widehat{\lambda}_y}(y) dy\right).$$

We use Corollary 6.4 to show the following intermediate statement.

Lemma 6.6. *Consider the setting of Lemma 6.2. For all $k \in \mathbb{N}_0$ and all $\mathbf{x} \in \Lambda'^k$ it holds that*

$$\lambda^{\mathbf{x}} \frac{Z_{(\Lambda')^c}(\lambda_{\mathbf{x}})}{Z_{\Lambda}(\lambda)} e^{-H(\mathbf{x})} \leq \exp\left(a \left|\Lambda'(r)\right| e^{br} \left(\lambda^{|\Lambda'|} + e^{\lambda^{|\Lambda'(r)|}}\right) e^{-b \cdot \text{dist}(\Lambda', \text{supp}(\lambda - \lambda'))}\right) \lambda^{\mathbf{x}} \frac{Z_{(\Lambda')^c}(\lambda'_{\mathbf{x}})}{Z_{\Lambda}(\lambda')} e^{-H(\mathbf{x})},$$

where $\Lambda'^{(r)} = \{y \in \Lambda \mid \text{dist}(y, \Lambda') < r\}$, and $Z_{\Lambda}(\lambda)$ and $Z_{\Lambda}(\lambda')$ are the partition functions on Λ for the potential ϕ .

Proof. Using Lemma 6.5 we have

$$\lambda^{\mathbf{x}} \frac{Z_{(\Lambda')^c}(\lambda_{\mathbf{x}})}{Z_{\Lambda}(\lambda)} e^{-H(\mathbf{x})} = \rho_{\lambda}(\mathbf{x}) \exp\left(-\int_{\Lambda'} \rho_{(\widehat{\lambda}_{\mathbf{x}})_y}(y) dy\right),$$

where we treat case that \mathbf{x} is the empty tuple by setting $\rho_{\lambda}(\mathbf{x}) = 1$. By Corollary 6.4, we have

$$\rho_{\lambda}(\mathbf{x}) \leq \left(1 + a \left|\Lambda'(r)\right| e^{br+\lambda^{|\Lambda'(r)|}} e^{-b \cdot \text{dist}(\Lambda', \text{supp}(\lambda - \lambda'))}\right) \rho_{\lambda'}(\mathbf{x})$$

$$\leq \exp\left(a\left|\Lambda^{(r)}\right|e^{br+\lambda\left|\Lambda^{(r)}\right|}e^{-b\cdot\text{dist}(\Lambda',\text{supp}(\boldsymbol{\lambda}-\boldsymbol{\lambda}'))}\right)\rho_{\boldsymbol{\lambda}'}(\mathbf{x}).$$

Moreover, observe that $\text{supp}\left(\widehat{(\boldsymbol{\lambda}_{\mathbf{x}})_y} - \widehat{(\boldsymbol{\lambda}'_{\mathbf{x}})_y}\right) \subseteq \text{supp}(\boldsymbol{\lambda} - \boldsymbol{\lambda}')$, $\widehat{(\boldsymbol{\lambda}_{\mathbf{x}})_y} \leq \boldsymbol{\lambda} < \lambda$ and $\widehat{(\boldsymbol{\lambda}'_{\mathbf{x}})_y} \leq \boldsymbol{\lambda}' < \lambda$ for all $y \in \Lambda'$. Thus, Lemma 6.2 implies

$$\begin{aligned} -\int_{\Lambda'} \rho_{\widehat{(\boldsymbol{\lambda}_{\mathbf{x}})_y}}(y) dy &\leq -\int_{\Lambda'} \rho_{\widehat{(\boldsymbol{\lambda}'_{\mathbf{x}})_y}}(y) dy + \int_{\Lambda'} \left| \rho_{\widehat{(\boldsymbol{\lambda}_{\mathbf{x}})_y}}(y) - \rho_{\widehat{(\boldsymbol{\lambda}'_{\mathbf{x}})_y}}(y) \right| dy \\ &\leq -\int_{\Lambda'} \rho_{\widehat{(\boldsymbol{\lambda}'_{\mathbf{x}})_y}}(y) dy + a\left|\Lambda^{(r)}\right|e^{br}e^{-b\cdot\text{dist}(\Lambda',\text{supp}(\boldsymbol{\lambda}-\boldsymbol{\lambda}'))} \cdot \int_{\Lambda'} \widehat{(\boldsymbol{\lambda}'_{\mathbf{x}})_y}(y) dy \\ &\leq -\int_{\Lambda'} \rho_{\widehat{(\boldsymbol{\lambda}'_{\mathbf{x}})_y}}(y) dy + a\left|\Lambda^{(r)}\right|e^{br}e^{-b\cdot\text{dist}(\Lambda',\text{supp}(\boldsymbol{\lambda}-\boldsymbol{\lambda}'))} \cdot \lambda\left|\Lambda'\right|. \end{aligned}$$

We conclude that

$$\begin{aligned} \lambda^{\mathbf{x}} \frac{Z_{(\Lambda')^c}(\boldsymbol{\lambda}_{\mathbf{x}})}{Z_{\Lambda}(\boldsymbol{\lambda})} e^{-H(\mathbf{x})} &\leq \exp\left(a\left|\Lambda^{(r)}\right|e^{br}\left(\lambda\left|\Lambda'\right| + e^{\lambda\left|\Lambda^{(r)}\right|}\right)e^{-b\cdot\text{dist}(\Lambda',\text{supp}(\boldsymbol{\lambda}-\boldsymbol{\lambda}'))}\right)\rho_{\boldsymbol{\lambda}'}(\mathbf{x}) \exp\left(-\int_{\Lambda'} \rho_{\widehat{(\boldsymbol{\lambda}'_{\mathbf{x}})_y}}(y) dy\right). \end{aligned}$$

Applying Lemma 6.5 again concludes the proof. \square

With Lemma 6.6, we can prove Lemma 6.1.

Proof of Lemma 6.1. We aim to apply Lemma 6.6. To this end, we start by writing the involved partition functions in terms of two new activity functions $\boldsymbol{\lambda}, \boldsymbol{\lambda}'$.

Set $\boldsymbol{\lambda} = \lambda_{\xi_1 \cup (\eta \cap \Lambda_{S \setminus \{\mathbf{v}\}})} \mathbb{1}_{\Lambda_B}$ and $\boldsymbol{\lambda}' = \lambda_{\xi_2 \cup (\eta \cap \Lambda_{S \setminus \{\mathbf{v}\}})} \mathbb{1}_{\Lambda_B}$, and observe that $\boldsymbol{\lambda}, \boldsymbol{\lambda}' \leq \lambda$. Let $\mathbf{x} \in \Lambda_{\mathbf{v}}^{|\eta \cap \Lambda_{\mathbf{v}}|}$ be any fixed tuple of points such that $\eta_{\mathbf{x}} = \eta \cap \Lambda_{\mathbf{v}}$ (i.e., any tuple containing exactly the points in $\eta \cap \Lambda_{\mathbf{v}}$). Note that

$$\lambda_{\xi_1 \cup (\eta \cap \Lambda_S)} \mathbb{1}_{\Lambda_B \setminus \{\mathbf{v}\}} = \boldsymbol{\lambda}_{\eta \cap \Lambda_{\mathbf{v}}} \mathbb{1}_{(\Lambda_{\mathbf{v}})^c} = \boldsymbol{\lambda}_{\mathbf{x}} \mathbb{1}_{(\Lambda_{\mathbf{v}})^c}$$

and analogously $\lambda_{\xi_2 \cup (\eta \cap \Lambda_S)} \mathbb{1}_{\Lambda_B \setminus \{\mathbf{v}\}} = \boldsymbol{\lambda}'_{\mathbf{x}} \mathbb{1}_{(\Lambda_{\mathbf{v}})^c}$. Thus, we obtain

$$\frac{Z_{\Lambda_B \setminus \{\mathbf{v}\}}(\lambda_{\xi_1 \cup (\eta \cap \Lambda_S)})}{Z_{\Lambda_B}(\lambda_{\xi_1 \cup (\eta \cap \Lambda_S \setminus \{\mathbf{v}\})})} \cdot \frac{Z_{\Lambda_B}(\lambda_{\xi_2 \cup (\eta \cap \Lambda_S \setminus \{\mathbf{v}\})})}{Z_{\Lambda_B \setminus \{\mathbf{v}\}}(\lambda_{\xi_2 \cup (\eta \cap \Lambda_S)})} = \frac{Z_{\Lambda}(\boldsymbol{\lambda}') Z_{(\Lambda_{\mathbf{v}})^c}(\boldsymbol{\lambda}_{\mathbf{x}})}{Z_{\Lambda}(\boldsymbol{\lambda}) Z_{(\Lambda_{\mathbf{v}})^c}(\boldsymbol{\lambda}'_{\mathbf{x}})} \quad (13)$$

We proceed by lower bounding the distance between $\text{supp}(\boldsymbol{\lambda} - \boldsymbol{\lambda}')$ and $\Lambda_{\mathbf{v}}$. Note that $\xi_1 \cup (\eta \cap \Lambda_{S \setminus \{\mathbf{v}\}})$ and $\xi_2 \cup (\eta \cap \Lambda_{S \setminus \{\mathbf{v}\}})$ agree on $\Lambda_{S \setminus \{\mathbf{v}\}}$ and can only disagree on Λ_H . By construction, it holds that $\text{dist}(\Lambda_H, \Lambda_{\mathbf{v}}) \geq \ell r$. As the potential range is bounded by r , it follows that $\text{dist}(\text{supp}(\boldsymbol{\lambda} - \boldsymbol{\lambda}'), \Lambda_{\mathbf{v}}) \geq (\ell - 1)r$.

Now, note that in particular $\boldsymbol{\lambda}$ and $\boldsymbol{\lambda}'$ agree on $\Lambda_{\mathbf{v}}$ and $\boldsymbol{\lambda}^{\mathbf{x}} = \boldsymbol{\lambda}'^{\mathbf{x}}$. We may assume $\lambda > 0$, since otherwise all involved partition functions are 1 and the statement holds trivially.

Since further η is feasible we have $\boldsymbol{\lambda}^{\mathbf{x}} e^{-H(\mathbf{x})} = \boldsymbol{\lambda}'^{\mathbf{x}} e^{-H(\mathbf{x})} > 0$. Thus, multiplying (13) with $\frac{\boldsymbol{\lambda}^{\mathbf{x}} e^{-H(\mathbf{x})}}{\boldsymbol{\lambda}'^{\mathbf{x}} e^{-H(\mathbf{x})}} = 1$ yields

$$\begin{aligned} \frac{Z_{\Lambda_B \setminus \{\mathbf{v}\}}(\lambda_{\xi_1 \cup (\eta \cap \Lambda_S)})}{Z_{\Lambda_B}(\lambda_{\xi_1 \cup (\eta \cap \Lambda_S \setminus \{\mathbf{v}\})})} \cdot \frac{Z_{\Lambda_B}(\lambda_{\xi_2 \cup (\eta \cap \Lambda_S \setminus \{\mathbf{v}\})})}{Z_{\Lambda_B \setminus \{\mathbf{v}\}}(\lambda_{\xi_2 \cup (\eta \cap \Lambda_S)})} &= \frac{\boldsymbol{\lambda}^{\mathbf{x}} Z_{\Lambda}(\boldsymbol{\lambda}') Z_{(\Lambda_{\mathbf{v}})^c}(\boldsymbol{\lambda}_{\mathbf{x}}) e^{-H(\mathbf{x})}}{\boldsymbol{\lambda}'^{\mathbf{x}} Z_{\Lambda}(\boldsymbol{\lambda}) Z_{(\Lambda_{\mathbf{v}})^c}(\boldsymbol{\lambda}'_{\mathbf{x}}) e^{-H(\mathbf{x})}} \\ &= \frac{\boldsymbol{\lambda}^{\mathbf{x}} \frac{Z_{(\Lambda_{\mathbf{v}})^c}(\boldsymbol{\lambda}'_{\mathbf{x}})}{Z_{\Lambda}(\boldsymbol{\lambda})} e^{-H(\mathbf{x})}}{\boldsymbol{\lambda}'^{\mathbf{x}} \frac{Z_{(\Lambda_{\mathbf{v}})^c}(\boldsymbol{\lambda}'_{\mathbf{x}})}{Z_{\Lambda}(\boldsymbol{\lambda}')} e^{-H(\mathbf{x})}}. \end{aligned}$$

Finally, by Lemma 6.6 we have the upper bound

$$\frac{Z_{\Lambda_B \setminus \{\mathbf{v}\}}(\lambda_{\xi_1 \cup (\eta \cap \Lambda_S)})}{Z_{\Lambda_B}(\lambda_{\xi_1 \cup (\eta \cap \Lambda_S \setminus \{\mathbf{v}\})})} \cdot \frac{Z_{\Lambda_B}(\lambda_{\xi_2 \cup (\eta \cap \Lambda_S \setminus \{\mathbf{v}\})})}{Z_{\Lambda_B \setminus \{\mathbf{v}\}}(\lambda_{\xi_2 \cup (\eta \cap \Lambda_S)})} \leq \exp\left(a 3^d r^d e^{2br} \left(\lambda r^d + e^{\lambda 3^d r^d}\right) e^{-br\ell}\right),$$

and applying the same reasoning after swapping the roles of $\boldsymbol{\lambda}$ and $\boldsymbol{\lambda}'$ results in the corresponding lower bound, which proves the claim. \square

7 Hard-sphere model

In this section we focus on the hard-sphere model. Recall that for an interaction range $r > 0$, the hard-sphere model is defined by the potential

$$\phi(x, y) = \begin{cases} \infty & \text{if } \text{dist}(x, y) < r \\ 0 & \text{otherwise} \end{cases}.$$

To simplify notation, define for every $\mathbf{x} \in \mathbb{R}^k$ and $\mathbf{y} \in \mathbb{R}^m$

$$\begin{aligned} D(\mathbf{x}) &= e^{-\sum_{\{i,j\} \in \binom{[k]}{2}} \phi(x_i, x_j)} = \prod_{\{i,j\} \in \binom{[k]}{2}} \mathbb{1}_{\text{dist}(x_i, x_j) \geq r} \text{ and} \\ D(\mathbf{x} \mid \mathbf{y}) &= e^{-\sum_{i \in [k]} \sum_{j \in [m]} \phi(x_i, y_j)} = \prod_{i \in [k]} \prod_{j \in [m]} \mathbb{1}_{\text{dist}(x_i, y_j) \geq r}. \end{aligned}$$

We extend this definition from tuples of points to finite point sets in the obvious way. This allows us to write

$$Z_{\Lambda'}(\lambda_{\eta}) = \sum_{k \geq 0} \frac{\lambda^k}{k!} \int_{\Lambda'^k} D(\mathbf{x}) D(\mathbf{x} \mid \eta) d\mathbf{x},$$

for all measurable $\Lambda' \subseteq \Lambda$ and all $\eta \in \mathcal{N}_{\Lambda}$.

7.1 Constructing the Bayes filter

We start by constructing a suitable Bayes filter correction for the hard-sphere model. The key ingredient will be computing such a correction by enumerating a finite set of boundary configurations that closely approximates all possible boundary conditions. This is made precise by the following lemma.

Lemma 7.1. *Let $S \subseteq \mathcal{V}$ be non-empty, $\mathbf{v} \in S$, $\eta \in \mathcal{N}_\Lambda$, $B = B(S, \mathbf{v}, \ell)$ and $H = (S \cup B)^c$. For all $\varepsilon > 0$, $\xi \in \mathcal{N}_{\Lambda_H}$ and*

$$\delta \leq \varepsilon \cdot \left(|\xi \cap \Lambda_{\partial B}| 2^d r^{d-1} d^{3/2} \lambda e^{\lambda(2\ell+1)d} r^d \right)^{-1}$$

there is some $\gamma \subseteq (\delta\mathbb{Z})^d \cap \Lambda_{H \cap \partial B}$ such that

$$e^{-\varepsilon} \cdot \frac{Z_{\Lambda_B \setminus \{\mathbf{u}_t\}}(\lambda_{\xi \cup (\eta \cap \Lambda_S)})}{Z_{\Lambda_B}(\lambda_{\xi \cup (\eta \cap \Lambda_S \setminus \{\mathbf{v}\})})} \leq \frac{Z_{\Lambda_B \setminus \{\mathbf{u}_t\}}(\lambda_{\gamma \cup (\eta \cap \Lambda_S)})}{Z_{\Lambda_B}(\lambda_{\gamma \cup (\eta \cap \Lambda_S \setminus \{\mathbf{v}\})})} \leq e^\varepsilon \cdot \frac{Z_{\Lambda_B \setminus \{\mathbf{u}_t\}}(\lambda_{\xi \cup (\eta \cap \Lambda_S)})}{Z_{\Lambda_B}(\lambda_{\xi \cup (\eta \cap \Lambda_S \setminus \{\mathbf{v}\})})}.$$

Proof. Fix some $\xi \in \mathcal{N}_{\Lambda_H}$, let Φ map every point in $\Lambda_{H \cap \partial B}$ to its closest point in $(\delta\mathbb{Z})^d \cap \Lambda_{H \cap \partial B}$ in ℓ_∞ -distance and set $\gamma = \{\Phi(x) \mid x \in \xi \cap \Lambda_{\partial B}\}$. We first prove that

$$e^{-\frac{\varepsilon}{2}} \cdot Z_{\Lambda_B \setminus \{\mathbf{u}_t\}}(\lambda_{\xi \cup (\eta \cap \Lambda_S)}) \leq Z_{\Lambda_B \setminus \{\mathbf{u}_t\}}(\lambda_{\gamma \cup (\eta \cap \Lambda_S)}) \leq e^{\frac{\varepsilon}{2}} \cdot Z_{\Lambda_B \setminus \{\mathbf{u}_t\}}(\lambda_{\xi \cup (\eta \cap \Lambda_S)}).$$

To this end, note that by Lemma C.2

$$Z_{\Lambda_B \setminus \{\mathbf{u}_t\}}(\lambda_{\xi \cup (\eta \cap \Lambda_S)}) = Z_{\Lambda_B \setminus \{\mathbf{u}_t\}}(\lambda_{(\xi \cap \Lambda_{\partial B}) \cup (\eta \cap \Lambda_S)}).$$

Thus, we have

$$\begin{aligned} & \left| Z_{\Lambda_B \setminus \{\mathbf{u}_t\}}(\lambda_{\xi \cup (\eta \cap \Lambda_S)}) - Z_{\Lambda_B \setminus \{\mathbf{u}_t\}}(\lambda_{\gamma \cup (\eta \cap \Lambda_S)}) \right| \\ & \leq \sum_{k \geq 1} \frac{\lambda^k}{k!} \int_{(\Lambda_B \setminus \{\mathbf{v}\})^k} D(\mathbf{x}) \cdot |D(\mathbf{x} \mid (\xi \cap \Lambda_{\partial B}) \cup (\eta \cap \Lambda_S)) - D(\mathbf{x} \mid \gamma \cup (\eta \cap \Lambda_S))| d\mathbf{x} \\ & \leq \sum_{k \geq 1} \frac{\lambda^k}{k!} \int_{(\Lambda_B \setminus \{\mathbf{v}\})^k} |D(\mathbf{x} \mid (\xi \cap \Lambda_{\partial B}) \cup (\eta \cap \Lambda_S)) - D(\mathbf{x} \mid \gamma \cup (\eta \cap \Lambda_S))| d\mathbf{x}. \end{aligned}$$

Next, observe that for every $\mathbf{x} \in (\Lambda_B \setminus \{\mathbf{v}\})^k$ it holds that $D(\mathbf{x} \mid (\xi \cap \Lambda_{\partial B}) \cup (\eta \cap \Lambda_S)) \neq D(\mathbf{x} \mid \gamma \cup (\eta \cap \Lambda_S))$ implies that there is some $i \in [k]$ and some $y \in \xi \cap \Lambda_{\partial B}$ such that either $\text{dist}(x_i, y) \geq r > \text{dist}(x_i, \Phi(y))$ or $\text{dist}(x_i, y) < r \leq \text{dist}(x_i, \Phi(y))$. Further, note that for every $y \in \Lambda_{H \cap \partial B}$ (and particular $y \in \xi \cap \Lambda_{\partial B}$) it holds that $\text{dist}(y, \Phi(y)) \leq \sqrt{d}\delta$. Using union bound and the fact that $D(\cdot \mid (\xi \cap \Lambda_{\partial B}) \cup (\eta \cap \Lambda_S))$ $D(\cdot \mid \gamma \cup (\eta \cap \Lambda_S))$ are symmetric functions, we obtain

$$\left| Z_{\Lambda_B \setminus \{\mathbf{u}_t\}}(\lambda_{\xi \cup (\eta \cap \Lambda_S)}) - Z_{\Lambda_B \setminus \{\mathbf{u}_t\}}(\lambda_{\gamma \cup (\eta \cap \Lambda_S)}) \right|$$

$$\begin{aligned}
&\leq |\xi \cap \Lambda_{\partial B}| \cdot \left[(r + \sqrt{d}\delta)^d - (r + \sqrt{d}\delta)^d \right] \cdot \sum_{k \geq 1} \frac{\lambda^k}{k!} k |\Lambda_{B \setminus \{\mathbf{v}\}}|^{k-1} \\
&\leq |\xi \cap \Lambda_{\partial B}| \cdot \left[(r + \sqrt{d}\delta)^d - (r + \sqrt{d}\delta)^d \right] \cdot \lambda e^{\lambda |\Lambda_{B \setminus \{\mathbf{v}\}}|}.
\end{aligned}$$

Elementary calculations yield

$$\left[(r + \sqrt{d}\delta)^d - (r + \sqrt{d}\delta)^d \right] \leq d^{3/2} \delta (r + \sqrt{d}\delta)^{d-1} \leq d^{3/2} \delta 2^{d-1} r^{d-1}.$$

Further, it holds that

$$|\Lambda_{B \setminus \{\mathbf{v}\}}| \leq |\Lambda_B| \leq (2\ell + 1)^d r^d.$$

Thus, for our choice of δ we obtain

$$\left| Z_{\Lambda_{B \setminus \{\mathbf{u}_t\}}}(\lambda_{\xi \cup (\eta \cap \Lambda_S)}) - Z_{\Lambda_{B \setminus \{\mathbf{u}_t\}}}(\lambda_{\gamma \cup (\eta \cap \Lambda_S)}) \right| \leq \frac{\varepsilon}{2},$$

and, since

$$\min \left\{ Z_{\Lambda_{B \setminus \{\mathbf{u}_t\}}}(\lambda_{\xi \cup (\eta \cap \Lambda_S)}), Z_{\Lambda_{B \setminus \{\mathbf{u}_t\}}}(\lambda_{\gamma \cup (\eta \cap \Lambda_S)}) \right\} \geq 1,$$

this proves the desired multiplicative bound.

It remains to show

$$e^{-\frac{\varepsilon}{2}} \cdot Z_{\Lambda_B}(\lambda_{\xi \cup (\eta \cap \Lambda_{S \setminus \{\mathbf{v}\}})}) \leq Z_{\Lambda_B}(\lambda_{\gamma \cup (\eta \cap \Lambda_{S \setminus \{\mathbf{v}\}})}) \leq e^{\frac{\varepsilon}{2}} \cdot Z_{\Lambda_B}(\lambda_{\xi \cup (\eta \cap \Lambda_{S \setminus \{\mathbf{v}\}})}),$$

which is done analogously, concluding the proof. \square

In particular, we obtain the following corollary.

Corollary 7.2. *Let $S \subseteq \mathcal{V}$ be non-empty, $\mathbf{v} \in S$, $\eta \in \mathcal{N}_\Lambda$, $B = B(S, \mathbf{v}, \ell)$ and $H = (S \cup B)^c$. For all $\varepsilon > 0$ and*

$$\delta \leq \varepsilon \cdot \left(4^d r^{-1} d^{(d+3)/2} (2\ell + 3)^d r^d \lambda e^{\lambda (2\ell+1)^d r^d} \right)^{-1}$$

it holds that

$$e^{-\varepsilon} \cdot \min_{\gamma \subseteq (\delta \mathcal{Z})^d \cap \Lambda_{H \cap \partial B}} \frac{Z_{\Lambda_{B \setminus \{\mathbf{v}\}}}(\lambda_{\gamma \cup (\eta \cap \Lambda_S)})}{Z_{\Lambda_B}(\lambda_{\gamma \cup (\eta \cap \Lambda_{S \setminus \{\mathbf{v}\}})})} \leq \inf_{\substack{\xi \in \mathcal{N}_{\Lambda_H} \\ \xi \cup (\eta \cap \Lambda_S) \text{ is feasible}}} \frac{Z_{\Lambda_{B \setminus \{\mathbf{v}\}}}(\lambda_{\xi \cup (\eta \cap \Lambda_S)})}{Z_{\Lambda_B}(\lambda_{\xi \cup (\eta \cap \Lambda_{S \setminus \{\mathbf{v}\}})})}.$$

Proof. Note that, if $\xi \in \mathcal{N}_{\Lambda_H}$ is such that $\xi \cup (\eta \cap \Lambda_S)$ is feasible, this implies in particular that $\xi' = \xi \cap \Lambda_{\partial B}$ is feasible. Thus, the claim follows from Lemma 7.1 by arguing that every feasible configuration $\xi' \in \mathcal{N}_{\Lambda_{H \cap \partial B}}$ satisfies $|\xi'| \leq \left(\frac{2\sqrt{d}}{r} \right)^d (2\ell + 3)^d r^d$. To see this, note that every point in ξ' blocks at least a volume of $\left(\frac{r}{2\sqrt{d}} \right)^d$, where no other point can be placed. Moreover, it holds that

$$|\Lambda_{H \cap \partial B}| \leq |\Lambda_{B \cup \partial B}| \leq (2\ell + 3)^d r^d,$$

which concludes the proof. \square

Corollary 7.2 allows us to replace the minimization over the uncountable set of boundary conditions \mathcal{N}_{Λ_H} by a minimization over the finite set $(\delta\mathbb{Z})^d \cap \Lambda_{H \cap \partial B}$. The second ingredient that we need for computing a suitable Bayes filter correction is a way to approximate the involved partition functions.

To this end, for every non-empty $S \subseteq \mathcal{V}$, $\eta \in \mathcal{N}_\Lambda$ and $\delta > 0$, define

$$\hat{Z}(S, \eta, \delta) = \sum_{\gamma \subseteq (\delta\mathbb{Z})^d \cap \Lambda_S} \lambda^{|\gamma|} \delta^{|\gamma|} \cdot D(\gamma) \cdot D(\gamma \mid \eta \cap \Lambda_{\partial S}). \quad (14)$$

The follow lemma justifies using $\hat{Z}(S, \eta, \delta)$ as an approximation for the hard-sphere partition function $Z_{\Lambda_S}(\lambda_\eta)$, given that δ is chosen sufficiently small.

Lemma 7.3. *Let $S \subseteq \mathcal{V}$ be non-empty and $\eta \in \mathcal{N}_\Lambda$. For all $\varepsilon > 0$ and*

$$\delta \leq \varepsilon \left[d^{3/2} 2^d \max\{r^{-1}, r^{d-1}\} \max\{\lambda, \lambda^2\} (|\eta \cap \Lambda_{\partial S}| + |\Lambda_{S \cup \partial S}|) e^{\lambda |\Lambda_{S \cup \partial S}|} \right]^{-1}$$

it holds that

$$e^{-\varepsilon} \cdot Z_{\Lambda_S}(\lambda_{\eta \cap (\Lambda_S)^c}) \leq \hat{Z}(S, \eta, \delta) \leq e^\varepsilon \cdot Z_{\Lambda_S}(\lambda_{\eta \cap (\Lambda_S)^c}),$$

where $\hat{Z}(S, \eta, \delta)$ is defined as in (14).

Proof. Define $\Lambda_{S, \delta} = \bigcup_{x \in (\delta\mathbb{Z})^d \cap \Lambda_S} \mathbb{B}_{\delta/2}^{(\infty)}(x)$, where $\mathbb{B}_{\delta/2}^{(\infty)}(x)$ is the closed $\delta/2$ -ball around x in infinity norm. Moreover, let Φ map every point in $\Lambda_{S, \delta}$ to its closest point in $(\delta\mathbb{Z})^d \cap \Lambda_S$ in ℓ_∞ -distance, breaking ties arbitrarily. With some abuse of notation, we extend Φ to tuples $\mathbf{x} \in \Lambda_{S, \delta}$ by setting $\Phi(\mathbf{x}) = (\Phi(x_1), \dots, \Phi(x_k))$. Now, note that

$$\begin{aligned} \hat{Z}(S, \eta, \delta) &= \sum_{k \geq 0} \frac{\lambda^k}{k!} \delta^k \sum_{\mathbf{x} \in (\delta\mathbb{Z})^{dk} \cap \Lambda_S^k} D(\mathbf{x}) \cdot D(\mathbf{x} \mid \eta \cap \Lambda_{\partial S}) \\ &= \sum_{k \geq 0} \frac{\lambda^k}{k!} \int_{\Lambda_{S, \delta}^k} D(\Phi(\mathbf{x})) \cdot D(\Phi(\mathbf{x}) \mid \eta \cap \Lambda_{\partial S}) d\mathbf{x}, \end{aligned}$$

where the first equality uses the fact $D(\mathbf{x}) = 0$ whenever \mathbf{x} contains the same point more than once.

We proceed by relating $\hat{Z}(S, \eta, \delta)$ to $Z_{\Lambda_S}(\lambda_{\eta \cap (\Lambda_S)^c})$ in two steps. First, we compare $Z_{\Lambda_S}(\lambda_{\eta \cap (\Lambda_S)^c})$ with $Z_{\Lambda_{S, \delta}}(\lambda_{\eta \cap \Lambda_{\partial S}})$, and then we compare $Z_{\Lambda_{S, \delta}}(\lambda_{\eta \cap \Lambda_{\partial S}})$ with $\hat{Z}(S, \eta, \delta)$ using the expression above.

For the first part, note that by Lemma C.2 it holds that $Z_{\Lambda_S}(\lambda_{\eta \cap (\Lambda_S)^c}) = Z_{\Lambda_S}(\lambda_{\eta \cap \Lambda_{\partial S}})$. Moreover, we have

$$\begin{aligned} Z_{\Lambda_S}(\lambda_{\eta \cap \Lambda_{\partial S}}) &\leq Z_{\Lambda_S \setminus \Lambda_{S, \delta}}(\lambda_{\eta \cap \Lambda_{\partial S}}) \cdot Z_{\Lambda_S \cap \Lambda_{S, \delta}}(\lambda_{\eta \cap \Lambda_{\partial S}}) \\ &\leq e^{\lambda |\Lambda_S \ominus \Lambda_{S, \delta}|} \cdot Z_{\Lambda_{S, \delta}}(\lambda_{\eta \cap \Lambda_{\partial S}}), \end{aligned}$$

where \ominus denotes the symmetric difference. Analogously, it holds that

$$Z_{\Lambda_{S,\delta}}(\lambda_{\eta \cap \Lambda_{\partial S}}) \leq e^{\lambda|\Lambda_S \ominus \Lambda_{S,\delta}|} \cdot Z_{\Lambda_S}(\lambda_{\eta \cap \Lambda_{\partial S}}).$$

Thus, if we show that $\lambda|\Lambda_S \ominus \Lambda_{S,\delta}| \leq \varepsilon/2$ for our choice of δ , then

$$e^{-\varepsilon/2} Z_{\Lambda_S}(\lambda_{\eta \cap (\Lambda_S)^c}) \leq Z_{\Lambda_{S,\delta}}(\lambda_{\eta \cap \Lambda_{\partial S}}) \leq e^{\varepsilon/2} Z_{\Lambda_S}(\lambda_{\eta \cap (\Lambda_S)^c}). \quad (15)$$

To this end, note that, if $x \in \Lambda_{S,\delta} \setminus \Lambda_S$, then $x \notin \Lambda_S$ but $\text{dist}(x, \Lambda_S) \leq \frac{\sqrt{d}\delta}{2}$. Similarly, if $x \in \Lambda_S \setminus \Lambda_{S,\delta}$, then $x \in \Lambda_S$ but $\text{dist}(x, (\Lambda_S)^c) \leq \frac{\sqrt{d}\delta}{2}$. Taking the union bound over $\mathbf{v} \in S$ yields

$$\begin{aligned} |\Lambda_S \ominus \Lambda_{S,\delta}| &\leq |S| \cdot \left[(r + \sqrt{d}\delta/2)^d - (r - \sqrt{d}\delta/2)^d \right] \\ &\leq |S| d^{3/2} \delta \cdot (r + \sqrt{d}\delta/2)^{d-1} \\ &\leq |\Lambda_S| d^{3/2} 2^{d-1} r^{-1} \delta. \end{aligned}$$

Thus, for $\delta \leq \varepsilon(\lambda|\Lambda_S|d^{3/2}2^{d-1}r^{-1})^{-1}$ the desired inequality is satisfied.

We proceed by relating $Z_{\Lambda_{S,\delta}}(\lambda_{\eta \cap \Lambda_{\partial S}})$ to $\hat{Z}(S, \eta, \delta)$. First, note that

$$\begin{aligned} &\left| Z_{\Lambda_{S,\delta}}(\lambda_{\eta \cap \Lambda_{\partial S}}) - \hat{Z}(S, \eta, \delta) \right| \\ &\leq \sum_{k \geq 0} \frac{\lambda^k}{k!} \int_{\Lambda_{S,\delta}^k} |D(\mathbf{x}) \cdot D(\mathbf{x} \mid \eta \cap \Lambda_{\partial S}) - D(\Phi(\mathbf{x})) \cdot D(\Phi(\mathbf{x}) \mid \eta \cap \Lambda_{\partial S})| d\mathbf{x} \\ &\leq \sum_{k \geq 2} \frac{\lambda^k}{k!} \int_{\Lambda_{S,\delta}^k} |D(\mathbf{x}) - D(\Phi(\mathbf{x}))| d\mathbf{x} + \sum_{k \geq 1} \frac{\lambda^k}{k!} \int_{\Lambda_{S,\delta}^k} |D(\mathbf{x} \mid \eta \cap \Lambda_{\partial S}) - D(\Phi(\mathbf{x}) \mid \eta \cap \Lambda_{\partial S})| d\mathbf{x}. \end{aligned}$$

We bound each of the terms in this sum separately. To this end, note that for any $\mathbf{x} \in \Lambda_{S,\delta}^k$ it holds that $D(\mathbf{x}) \neq D(\Phi(\mathbf{x}))$ implies that there are $i < j$ such that either $\text{dist}(x_i, x_j) < r \leq \text{dist}(\Phi(x_i), \Phi(x_j))$ or $\text{dist}(x_i, x_j) \geq r > \text{dist}(\Phi(x_i), \Phi(x_j))$. Since $\text{dist}(x, \Phi(x)) \leq \frac{\sqrt{d}\delta}{2}$, applying union bound over $1 \leq i < j \leq k$ yields

$$\begin{aligned} \sum_{k \geq 2} \frac{\lambda^k}{k!} \int_{\Lambda_{S,\delta}^k} |D(\mathbf{x}) - D(\Phi(\mathbf{x}))| d\mathbf{x} &\leq \left[(r + \sqrt{d}\delta)^d - (r - \sqrt{d}\delta)^d \right] \sum_{k \geq 2} \frac{\lambda^k}{k!} k(k-1) |\Lambda_{S,\delta}|^{k-1} \\ &\leq d^{3/2} \delta \cdot (r + \sqrt{d}\delta)^{d-1} \lambda^2 |\Lambda_{S,\delta}| e^{\lambda|\Lambda_{S,\delta}|} \\ &\leq d^{3/2} \delta 2^{d-1} r^{d-1} \lambda^2 |\Lambda_{S,\delta}| e^{\lambda|\Lambda_{S,\delta}|}. \end{aligned}$$

Similarly, we have

$$\sum_{k \geq 1} \frac{\lambda^k}{k!} \int_{\Lambda_{S,\delta}^k} |D(\mathbf{x} \mid \eta \cap \Lambda_{\partial S}) - D(\Phi(\mathbf{x}) \mid \eta \cap \Lambda_{\partial S})| d\mathbf{x} \leq d^{3/2} \delta 2^{d-1} r^{d-1} \lambda |\eta \cap \Lambda_{\partial S}| e^{\lambda|\Lambda_{S,\delta}|}.$$

Combining both and noting that $\Lambda_{S,\delta} \subseteq \bigcup_{x \in \Lambda_S} \mathbb{B}_{\delta/2}^{(\infty)}(x) \subseteq \Lambda_{S \cup \partial S}$ yields

$$\left| Z_{\Lambda_{S,\delta}}(\lambda_{\eta \cap \Lambda_{\partial S}}) - \hat{Z}(S, \eta, \delta) \right| \leq d^{3/2} \delta 2^{d-1} r^{d-1} \max\{\lambda, \lambda^2\} (|\eta \cap \Lambda_{\partial S}| + |\Lambda_{S \cup \partial S}|) e^{\lambda |\Lambda_{S \cup \partial S}|}.$$

For $\delta \leq \varepsilon (d^{3/2} 2^d r^{d-1} \max\{\lambda, \lambda^2\} (|\eta \cap \Lambda_{\partial S}| + |\Lambda_{S \cup \partial S}|) e^{\lambda |\Lambda_{S \cup \partial S}|})^{-1}$ this gives

$$\left| Z_{\Lambda_{S,\delta}}(\lambda_{\eta \cap \Lambda_{\partial S}}) - \hat{Z}(S, \eta, \delta) \right| \leq \varepsilon/2,$$

and, since $Z_{\Lambda_{S,\delta}}(\lambda_{\eta \cap \Lambda_{\partial S}}) \geq 1$ and $\hat{Z}(S, \eta, \delta) \geq 1$,

$$e^{-\varepsilon/2} \cdot Z_{\Lambda_{S,\delta}}(\lambda_{\eta \cap \Lambda_{\partial S}}) \leq \hat{Z}(S, \eta, \delta) \leq e^{\varepsilon/2} \cdot Z_{\Lambda_{S,\delta}}(\lambda_{\eta \cap \Lambda_{\partial S}}).$$

Combining this with (15) concludes the proof. \square

We now combine Corollary 7.2 and Lemma 7.3 to obtain our Bayes filter correction for the hard-sphere model.

Lemma 7.4. *For $\varepsilon > 0$, non-empty $S \subseteq \mathcal{V}$, $\mathbf{v} \in S$ and feasible $\eta \in \mathcal{N}_\Lambda$ set*

$$\begin{aligned} \delta_1 &= \delta_1(\varepsilon) := \frac{\varepsilon}{2} \cdot \left(4^d r^{-1} d^{(d+3)/2} (2\ell + 3)^{d_r d} \lambda e^{\lambda(2\ell+1)d_r d} \right)^{-1} \text{ and} \\ \delta_2 &= \delta_2(\varepsilon) := \frac{\varepsilon}{4} \cdot \left(d^{3/2} 2^d \max\{r^{-1}, r^{d-1}\} \max\{\lambda, \lambda^2\} m e^{\lambda(2\ell+3)d_r d} \right)^{-1}, \end{aligned}$$

where $m = 2^d (2\ell + 3)^{d_r d} (\delta_1^{-d} + d^{d/2} r^{-d}) + (2\ell + 3)^{d_r d}$, and define

$$C_\varepsilon(S, \mathbf{v}, \eta) := e^{-\varepsilon} \cdot \min_{\gamma \subseteq (\delta_1 \mathbb{Z})^d \cap \Lambda_{H \cap \partial B}} \frac{\hat{Z}(B \setminus \{\mathbf{v}\}, \gamma \cup (\eta \cap \Lambda_S), \delta_2)}{\hat{Z}(B, \gamma \cup (\eta \cap \Lambda_{S \setminus \{\mathbf{v}\}}), \delta_2)},$$

where $B = B(S, \mathbf{v}, \ell)$ and $H = (S \cup B)^c$. Then $C_\varepsilon(S, \mathbf{v}, \eta)$ is a Bayes filter correction as in Definition 4.1.

Proof. We start by arguing that C_ε is a Bayes filter correction. For the measurability, note that for every fixed non-empty $S \subseteq \mathcal{V}$ and $\mathbf{v} \in S$ it holds that $C_\varepsilon(S, \mathbf{v}, \cdot)$ is a minimum of a finite set of \mathfrak{R}_{Λ_S} -measurable functions. In particular, it holds that $C_\varepsilon(S, \mathbf{v}, \eta) = C_\varepsilon(S, \mathbf{v}, \eta')$ for every two configurations $\eta, \eta' \in \mathcal{N}_\Lambda$ that agree on Λ_S .

Next, we argue that for all feasible $\eta \in \mathcal{N}_\Lambda$ it holds that $C_\varepsilon(S, \mathbf{v}, \eta)$ is bounded away from 0 and

$$C_\varepsilon(S, \mathbf{v}, \eta) \leq \inf_{\substack{\xi \in \mathcal{N}_{\Lambda_H} \\ \xi \cup (\eta \cap \Lambda_S) \text{ is feasible}}} \frac{Z_{\Lambda_{B \setminus \{\mathbf{v}\}}}(\lambda_{\xi \cup (\eta \cap \Lambda_S)})}{Z_{\Lambda_B}(\lambda_{\xi \cup (\eta \cap \Lambda_{S \setminus \{\mathbf{v}\}})})}.$$

For the lower bound, note that for all $\gamma \subseteq (\delta_1 \mathbb{Z})^d \cap \Lambda_{H \cap \partial B}$ it holds that

$$\frac{\hat{Z}(B \setminus \{\mathbf{v}\}, \gamma \cup (\eta \cap \Lambda_S), \delta_2)}{\hat{Z}(B, \gamma \cup (\eta \cap \Lambda_{S \setminus \{\mathbf{v}\}}), \delta_2)} \geq (1 + \lambda \delta_2)^{-|(\delta_2 \mathbb{Z})^d \cap \Lambda_B|} > 0$$

independent of η . For the upper bound, we start by observing that, for our choice of δ_1 , Corollary 7.2 yields

$$e^{-\varepsilon/2} \cdot \min_{\gamma \subseteq (\delta_1 \mathbb{Z})^d \cap \Lambda_{H \cap \partial B}} \frac{Z_{\Lambda_B \setminus \{\mathbf{v}\}}(\lambda_{\gamma \cup (\eta \cap \Lambda_S)})}{Z_{\Lambda_B}(\lambda_{\gamma \cup (\eta \cap \Lambda_S \setminus \{\mathbf{v}\})})} \leq \inf_{\substack{\xi \in \mathcal{N}_{\Lambda_H} \\ \xi \cup (\eta \cap \Lambda_S) \text{ is feasible}}} \frac{Z_{\Lambda_B \setminus \{\mathbf{v}\}}(\lambda_{\xi \cup (\eta \cap \Lambda_S)})}{Z_{\Lambda_B}(\lambda_{\xi \cup (\eta \cap \Lambda_S \setminus \{\mathbf{v}\})})}.$$

Note that $|\Lambda_B \setminus \{\mathbf{v}\}| \leq |\Lambda_B| \leq (2\ell + 1)^{d_r d}$ and $|\Lambda_B \setminus \{\mathbf{v}\} \cup \partial(B \setminus \{\mathbf{v}\})| \leq |\Lambda_{B \cup \partial B}| \leq (2\ell + 3)^{d_r d}$. Moreover, we have the crude bound $|(\delta_1 \mathbb{Z})^d \cap \Lambda_{H \cap \partial B}| \leq 2^d \delta_1^{-d} (2\ell + 3)^{d_r d}$, and, for every feasible $\eta \in \mathcal{N}_\Lambda$, it holds that $|\eta \cap \Lambda_{S \cap \partial(B \setminus \{\mathbf{v}\})}| \leq 2^d d^{d/2} r^{-d} (2\ell + 3)^{d_r d}$. Therefore, for all $\gamma \subseteq (\delta_1 \mathbb{Z})^d \cap \Lambda_{H \cap \partial B}$ we have $|(\gamma \cup (\eta \cap \Lambda_S)) \cap \Lambda_{\partial(B \setminus \{\mathbf{v}\})}| \leq 2^d (2\ell + 3)^{d_r d} (\delta_1^{-d} + d^{d/2} r^{-d})$. Analogously, it holds that $|(\gamma \cup (\eta \cap \Lambda_S \setminus \{\mathbf{v}\})) \cap \Lambda_{\partial B}| \leq 2^d (2\ell + 3)^{d_r d} (\delta_1^{-d} + d^{d/2} r^{-d})$. Thus, Lemma 7.3 yields for our choice of δ_2

$$e^{-\varepsilon/2} \cdot \frac{Z_{\Lambda_B \setminus \{\mathbf{v}\}}(\lambda_{\gamma \cup (\eta \cap \Lambda_S)})}{Z_{\Lambda_B}(\lambda_{\gamma \cup (\eta \cap \Lambda_S \setminus \{\mathbf{v}\})})} \leq \frac{\hat{Z}(B \setminus \{\mathbf{v}\}, \gamma \cup (\eta \cap \Lambda_S), \delta_2)}{\hat{Z}(B, \gamma \cup (\eta \cap \Lambda_S \setminus \{\mathbf{v}\}), \delta_2)} \leq e^{\varepsilon/2} \cdot \frac{Z_{\Lambda_B \setminus \{\mathbf{v}\}}(\lambda_{\gamma \cup (\eta \cap \Lambda_S)})}{Z_{\Lambda_B}(\lambda_{\gamma \cup (\eta \cap \Lambda_S \setminus \{\mathbf{v}\})})}.$$

In particular, this proves

$$e^{-\varepsilon} \cdot \min_{\gamma \subseteq (\delta_1 \mathbb{Z})^d \cap \Lambda_{H \cap \partial B}} \frac{\hat{Z}(B \setminus \{\mathbf{v}\}, \gamma \cup (\eta \cap \Lambda_S), \delta_2)}{\hat{Z}(B, \gamma \cup (\eta \cap \Lambda_S \setminus \{\mathbf{v}\}), \delta_2)} \leq \inf_{\substack{\xi \in \mathcal{N}_{\Lambda_H} \\ \xi \cup (\eta \cap \Lambda_S) \text{ is feasible}}} \frac{Z_{\Lambda_B \setminus \{\mathbf{v}\}}(\lambda_{\xi \cup (\eta \cap \Lambda_S)})}{Z_{\Lambda_B}(\lambda_{\xi \cup (\eta \cap \Lambda_S \setminus \{\mathbf{v}\})})},$$

implying that C_ε is a Bayes filter correction. \square

7.2 Efficiency of the algorithm

We now argue that under the assumption of strong spatial mixing we can use C_ε to obtain an efficient implementation of Algorithm 1. Our argument will consist of two steps. First, we discuss how to implement each step of the algorithm efficiently. In particular, we argue that we can efficiently update the configuration (line 7), and that we can efficiently sample a Bayes filter based on $C_\varepsilon(\cdot)$ (line 5). For the latter, we make use of a Bernoulli factory to circumvent the lack of an algorithm for exact computation of partition functions. In the second part, we focus on the overall number of iterations of the algorithm. This is where the assumption of strong spatial mixing comes into play to ensure that the success probability of our Bayes filter is sufficiently large, which implies rapid termination of the algorithm.

We start with discussing the running time of each iteration of Algorithm 1. For updating the configuration, we will use a rejection sampling method which, as long as the updated region Λ_B has constant volume, will be efficient enough for our setting. Since apply the same argument for more general repulsive potentials, the following lemma is stated in this general setting.

Lemma 7.5. *Let $B \subseteq \mathcal{V}$ and $\eta \in \mathcal{N}_\Lambda$. For any repulsive finite-range potential ϕ we can sample from $\mu_{\lambda_{\eta \cap (\Lambda_B)^c} \mathbb{1}_{\Lambda_B}}[\Lambda_B]$ in expected time $(\lambda|\Lambda_B| + |\eta \cap \Lambda_{\partial B}|) \cdot \lambda|\Lambda_B|e^{\lambda|\Lambda_B|}$.*

Proof. Let $P_{\Lambda_B, \lambda}$ denote a Poisson point process on Λ_B with intensity λ . We consider the rejection sampling algorithm given in Algorithm 2.

Algorithm 2: Sample from $\mu_{\lambda_{\eta \cap (\Lambda_B)^c} \mathbb{1}_{\Lambda_B}}[\Lambda_B]$

```

1 repeat
2   Draw  $Y \sim P_{\Lambda_B, \lambda}$ 
3   Compute  $w = e^{-\sum_{\{x, y\} \in \binom{Y}{2}} \phi(x, y)} \cdot e^{-\sum_{x \in \eta \cap \Lambda_{\partial B}, y \in Y} \phi(x, y)}$ 
4   Draw  $W \sim \text{Ber}(w)$ 
5 until  $W = 1$ ;
6 return  $Y$ 

```

To prove that this rejection sampling method produces the correct output distribution, it suffices to argue that w is proportional to the density of $\mu_{\lambda_{\eta \cap (\Lambda_B)^c} \mathbb{1}_{\Lambda_B}}[\Lambda_B]$ with respect to $P_{\Lambda_B, \lambda}$ for $P_{\Lambda_B, \lambda}$ -almost all $\xi \in \mathcal{N}_{\Lambda_B}$. This is true since for all $\xi \in \mathcal{N}_{\Lambda_B}$

$$\begin{aligned} \sum_{x \in \eta \cap (\Lambda_B)^c, y \in \xi} \phi(x, y) &= \sum_{x \in \eta \cap \Lambda_{\partial B}, y \in \xi} \phi(x, y) + \sum_{x \in \eta \cap (\Lambda_{B \cup \partial B})^c, y \in \xi} \phi(x, y) \\ &= \sum_{x \in \eta \cap \Lambda_{\partial B}, y \in \xi} \phi(x, y), \end{aligned}$$

where the last equality follows from the fact that $\text{dist}(\Lambda_B, (\Lambda_{B \cup \partial B})^c) \geq r$ and therefore $\phi(x, y) = 0$ for all $x \in (\Lambda_{B \cup \partial B})^c, y \in \Lambda_B$.

We proceed by using Wald's identity as given in Lemma B.1 to bound the expected running time of the procedure above. To this end, let $(Y_n)_{n \in \mathbb{N}}$ be a sequence of independent samples from $P_{\Lambda_B, \lambda}$. Assume the algorithm draws $Y = Y_n$ at iteration $n \in \mathbb{N}$. Let S_n denote the running time of the rejection sampler in iteration n and let N be the (random) number of iterations until the algorithm terminates. We aim for bounding $\mathbb{E}\left[\sum_{n=1}^N S_n\right]$.

First, observe that S_n is dominated by the time for computing w , implying $S_n \leq |Y_n|^2 + |Y_n| \cdot |\eta \cap \Lambda_B|$. Since further $|Y_n|$ follows a Poisson distribution with parameter $\lambda|\Lambda_B|$, we have $\mathbb{E}[S_n] \leq \lambda^2|\Lambda_B|^2 + \lambda|\Lambda_B| \cdot |\eta \cap \Lambda_{\partial B}| = (\lambda|\Lambda_B| + |\eta \cap \Lambda_{\partial B}|) \cdot \lambda|\Lambda_B|$.

Moreover, observe that the random variable $\mathbb{1}_{N \geq n}$ only depends on $(Y_i)_{i \leq n-1}$, whereas S_n only depends on Y_n . Therefore, S_n and $\mathbb{1}_{N \geq n}$ are independent and $\mathbb{E}[S_n \mathbb{1}_{N \geq n}] = \mathbb{E}[S_n] \mathbb{E}[\mathbb{1}_{N \geq n}]$.

Applying Lemma B.1 yields $\mathbb{E}\left[\sum_{n=1}^N S_n\right] \leq (\lambda|\Lambda_B| + |\eta \cap \Lambda_{\partial B}|) \cdot \lambda|\Lambda_B| \cdot \mathbb{E}[N]$. To obtain a bound $\mathbb{E}[N]$ on, observe that the algorithm always terminates if $Y = \emptyset$, which happens in every iteration independently with a probability of $e^{-\lambda|\Lambda_B|}$. Thus, N is dominated by a geometric random variable with success probability $e^{-\lambda|\Lambda_B|}$ and $\mathbb{E}[N] \leq e^{\lambda|\Lambda_B|}$, which concludes the proof. \square

Note that, in the case of the hard-sphere model with $r > 0$, if $\eta \in \mathcal{N}_\Lambda$ is feasible, then $|\eta \cap \Lambda_{\partial B}|$ is a linear function of the volume $|\Lambda_{\partial B}|$.

We proceed with bounding the running time for sampling the Bayes filter for the hard-sphere model in each step. To this end, we start with the following observation.

Observation 7.6. *Consider the setting of Lemma 7.4. The required running time for computing $C_\varepsilon(S, \mathbf{v}, \eta)$ does only depend on $\varepsilon, \ell, r, \lambda$ and d .*

This follows directly from enumerating all subsets $\gamma \subseteq (\delta_1 \mathbb{Z})^d \cap \Lambda_{H \cap \partial B}$ and brute-force computation of $\hat{Z}(B \setminus \{\mathbf{v}\}, \gamma \cup (\eta \cap \Lambda_S), \delta_2)$ and $\hat{Z}(B, \gamma \cup (\eta \cap \Lambda_{S \setminus \mathbf{v}}), \delta_2)$, where δ_1, δ_2 are as in Lemma 7.4.

In fact, we will not use C_ε directly for our Bayes filter, but rather a slightly scaled version $e^{-\varepsilon} C_\varepsilon$, which is again a Bayes filter correction. The slack due to the additional scaling allows us to efficiently sample the Bayes filter by using a Bernoulli factory, as we argue in the next lemma.

Lemma 7.7. *Let $S \subseteq \mathcal{V}$ be non-empty, $\mathbf{v} \in S$ and $\eta \in \mathcal{N}_\Lambda$ be feasible, and set $B = B(S, \mathbf{v}, \ell)$. For all $\varepsilon > 0$ we can sample a Bernoulli random variable with success probability*

$$e^{-\varepsilon} C_\varepsilon(S, \mathbf{v}, \eta) \cdot \frac{Z_{\Lambda_B}(\lambda_{\eta \cap \Lambda_{\partial B}})}{Z_{\Lambda_{B \setminus \{\mathbf{v}\}}}(\lambda_{\eta \cap \Lambda_{\partial B \cup \{\mathbf{v}\}}})}$$

with expected running time only depending on $\varepsilon, \ell, r, \lambda$ and d .

Proof. Our goal is to use a Bernoulli factory of the form $\frac{p}{q}$ to perform this task. To bring the desired success probability into such a form, note that

$$\begin{aligned} Z_{\Lambda_B}(\lambda_{\eta \cap \Lambda_{\partial B}}) &= \frac{1}{\mu_{\lambda_{\eta \cap \Lambda_{\partial B}}} \mathbb{1}_{\Lambda_B}(\{\emptyset\})} \\ Z_{\Lambda_{B \setminus \{\mathbf{v}\}}}(\lambda_{\eta \cap \Lambda_{\partial B \cup \{\mathbf{v}\}}}) &= \frac{1}{\mu_{\lambda_{\eta \cap \Lambda_{\partial B \cup \{\mathbf{v}\}}}} \mathbb{1}_{\Lambda_{B \setminus \{\mathbf{v}\}}}(\{\emptyset\})}. \end{aligned}$$

Moreover, note that $0 \leq e^{-\varepsilon} C_\varepsilon(S, \mathbf{v}, \eta) \leq 1$. Thus, by setting $p = e^{-\varepsilon} C_\varepsilon(S, \mathbf{v}, \eta) \cdot \mu_{\lambda_{\eta \cap \Lambda_{\partial B \cup \{\mathbf{v}\}}}} \mathbb{1}_{\Lambda_{B \setminus \{\mathbf{v}\}}}(\{\emptyset\})$ and $q = \mu_{\lambda_{\eta \cap \Lambda_{\partial B}}} \mathbb{1}_{\Lambda_B}(\{\emptyset\})$ we have $p \in [0, 1], q \in [0, 1]$ and

$$e^{-\varepsilon} C_\varepsilon(S, \mathbf{v}, \eta) \cdot \frac{Z_{\Lambda_B}(\lambda_{\eta \cap \Lambda_{\partial B}})}{Z_{\Lambda_{B \setminus \{\mathbf{v}\}}}(\lambda_{\eta \cap \Lambda_{\partial B \cup \{\mathbf{v}\}}})} = \frac{p}{q}.$$

We are now going to use Lemma 3.1 to prove that we can obtain a sample from $\text{Ber}\left(\frac{p}{q}\right)$ within the desired expected running time. To this end, we need to provide a positive lower bound on $q - p$ and we need an efficient way for generating independent samples from $\text{Ber}(q)$ and $\text{Ber}(p)$.

For the lower bound, note that by Lemma C.2 $Z_{\Lambda_B}(\lambda_{\eta \cap \Lambda_{\partial B}}) = Z_{\Lambda_B}(\lambda_{\eta \cap (\Lambda_B)^c})$ and $Z_{\Lambda_{B \setminus \{\mathbf{v}\}}}(\lambda_{\eta \cap \Lambda_{\partial B \cup \{\mathbf{v}\}}}) = Z_{\Lambda_{B \setminus \{\mathbf{v}\}}}(\lambda_{\eta \cap (\Lambda_{B \setminus \{\mathbf{v}\}})^c})$. Moreover, since $C_\varepsilon(S, \mathbf{v}, \eta)$ is a Bayes filter correction by Lemma 7.4 and η is feasible, we have

$$C_\varepsilon(S, \mathbf{v}, \eta) \cdot \frac{Z_{\Lambda_B}(\lambda_{\eta \cap \Lambda_{\partial B}})}{Z_{\Lambda_{B \setminus \{\mathbf{v}\}}}(\lambda_{\eta \cap \Lambda_{\partial B \cup \{\mathbf{v}\}}})} \leq 1.$$

Consequently, $\frac{p}{q} \leq e^{-\varepsilon}$ and

$$q - p \geq (1 - e^{-\varepsilon})q = (1 - e^{-\varepsilon}) \cdot (Z_{\Lambda_B}(\lambda_{\eta \cap \Lambda_{\partial B}}))^{-1} \geq (1 - e^{-\varepsilon})e^{-\lambda|\Lambda_B|}.$$

Using the upper bound $|\Lambda_B| \leq (2\ell + 1)^d r^d$ yields $q - p \geq (1 - e^{-\varepsilon})e^{-\lambda(2\ell+1)^d r^d}$.

We proceed by arguing that we can obtain an oracle for $\text{Ber}(p)$ and $\text{Ber}(q)$ as required by Lemma 3.1. Firstly, note that by Observation 7.6 we can compute $C_\varepsilon(S, \mathbf{v}, \eta)$ with running time only depending on $\varepsilon, \ell, r, \lambda$ and d . After computing $C_\varepsilon(S, \mathbf{v}, \eta)$, each independent sample from $\text{Ber}(e^{-\varepsilon}C_\varepsilon(S, \mathbf{v}, \eta))$ can be obtained in constant time. Thus, it remains to argue that we can efficiently sample independent Bernoulli random variables with success probabilities $\mu_{\lambda_{\eta \cap \Lambda_{\partial B \cup \{\mathbf{v}\}}} \mathbb{1}_{\Lambda_{B \setminus \{\mathbf{v}\}}}}(\{\emptyset\})$ and $\mu_{\lambda_{\eta \cap \Lambda_{\partial B}} \mathbb{1}_{\Lambda_B}}(\{\emptyset\})$. To this end, note that for $Y \sim \mu_{\lambda_{\eta \cap \Lambda_{\partial B}} \mathbb{1}_{\Lambda_B}}$ it holds that $Y \cap (\Lambda_B)^c = \emptyset$ almost surely. Thus, it holds that

$$\mu_{\lambda_{\eta \cap \Lambda_{\partial B}} \mathbb{1}_{\Lambda_B}}(\{\emptyset\}) = \mu_{\lambda_{\eta \cap \Lambda_{\partial B}} \mathbb{1}_{\Lambda_B}}[\Lambda_B](\{\emptyset\}).$$

Observe that for $Y \sim \mu_{\lambda_{\eta \cap \Lambda_{\partial B}} \mathbb{1}_{\Lambda_B}}[\Lambda_B]$ it holds that $\mathbb{1}_{Y=\emptyset}$ is a Bernoulli random variable with success probability $\mu_{\lambda_{\eta \cap \Lambda_{\partial B}} \mathbb{1}_{\Lambda_B}}(\{\emptyset\})$. By Lemma 7.5, we can obtain independent samples from $\mu_{\lambda_{\eta \cap \Lambda_{\partial B}} \mathbb{1}_{\Lambda_B}}[\Lambda_B]$, each in expected time at most $(\lambda|\Lambda_B| + |\eta \cap \Lambda_{\partial B}|) \cdot \lambda|\Lambda_B|e^{\lambda|\Lambda_B|}$. Note that $|\Lambda_B| \leq (2\ell + 1)^d r^d$ and that, for feasible $\eta \in \mathcal{N}_\Lambda$, $|\eta \cap \Lambda_{\partial B}| \leq (2\sqrt{d}/r)^d |\Lambda_{\partial B}| \leq (2\sqrt{d}/r)^d \cdot (2\ell + 3)^d r^d$. Therefore, the expected running time for obtaining independent Bernoulli samples with success probability $\mu_{\lambda_{\eta \cap \Lambda_{\partial B}} \mathbb{1}_{\Lambda_B}}(\{\emptyset\})$ is bounded by some function of $\varepsilon, \ell, r, \lambda$ and d .

Treating $\mu_{\lambda_{\eta \cap \Lambda_{\partial B \cup \{\mathbf{v}\}}} \mathbb{1}_{\Lambda_{B \setminus \{\mathbf{v}\}}}}(\{\emptyset\})$ analogously and applying Lemma 3.1 now proves our claim. \square

We conclude the following bound on the running time of each iteration.

Corollary 7.8. *Suppose we run Algorithm 1 on a hard-sphere model with $C(\cdot) = e^{-\varepsilon}C_\varepsilon(\cdot)$ as Bayes filter correction in line 5 for some $\varepsilon > 0$, and let R_t denote the running time of iteration $t \in \mathbb{N}$. Then, for all $t \in \mathbb{N}$, $\mathbb{E}[R_t \mid X_{t-1}, \mathcal{U}_{t-1}, \mathbf{u}_{t-1}]$ is almost surely bounded by some function of $\varepsilon, \ell, r, \lambda$ and d .*

Proof. Set $B = B(\mathcal{U}_{t-1}, \mathbf{u}_{t-1}, r)$, and note that the bulk of the running time in each iteration of Algorithm 1 is due to sampling the Bayes filter in line 5 and updating the point configuration on Λ_B in line 7.

For line 5, note that X_{t-1} is almost surely feasible by Lemma 4.3. Thus, Lemma 7.7 yields that the expected time for sampling the Bayes filter, conditioned on $X_{t-1}, \mathcal{U}_{t-1}$ and \mathbf{u}_{t-1} , almost surely bounded by some function of $\varepsilon, \ell, r, \lambda$ and d .

For line 7, we can use Lemma 7.5 to bound the expected time for sampling from $\mu_{\lambda_{X_{t-1} \cap (\Lambda_B)^c} \mathbb{1}_{\Lambda_B}}[\Lambda_B]$ is bounded by $(\lambda|\Lambda_B| + |X_{t-1} \cap \Lambda_{\partial B}|) \cdot \lambda|\Lambda_B|e^{\lambda|\Lambda_B|}$. Note that $|\Lambda_B| \leq (2\ell + 1)^{d_r d}$ and, if X_{t-1} is feasible, $|X_{t-1} \cap \Lambda_{\partial B}| \leq (2\sqrt{d}/r)^d |\Lambda_{\partial B}| \leq (2\sqrt{d}/r)^d \cdot (2\ell + 3)^{d_r d}$. Since X_{t-1} is indeed almost surely feasible, the expected running for line 7, conditioned on $X_{t-1}, \mathcal{U}_{t-1}$ and \mathbf{u}_{t-1} , is almost surely bounded by some function of $\varepsilon, \ell, r, \lambda$ and d as well, which concludes the proof. \square

We proceed by bounding the expected number of iterations of Algorithm 1, running on a hard-sphere model. To this end, we start with the following lower bound on the success probability of the Bayes filter with correction $e^{-\varepsilon} C_\varepsilon(\cdot)$ for a particular choice of ε .

Lemma 7.9. *Consider a hard-sphere model that exhibits (a, b) -strong spatial mixing up to λ . Then there are constants a', b' , only depending on a, b, r, λ and d , such that for all non-empty $S \subseteq \mathcal{V}$, $\mathbf{v} \in S$ and feasible $\eta \in \mathcal{N}_\Lambda$ it holds that*

$$e^{-e^{-\ell}} C_{e^{-\ell}}(S, \mathbf{v}, \eta) \cdot \frac{Z_{\Lambda_B}(\lambda_{\eta \cap \Lambda_{\partial B}})}{Z_{\Lambda_{B \setminus \{\mathbf{v}\}}}(\lambda_{\eta \cap \Lambda_{\partial B \cup \{\mathbf{v}\}}})} \geq 1 - a' e^{-b' \ell}.$$

Proof. Set $\delta_1 = \delta_1(e^{-\ell})$ and $\delta_2 = \delta_2(e^{-\ell})$ as defined in Lemma 7.4. Note that by Lemma C.2 we have

$$Z_{\Lambda_B}(\lambda_{\eta \cap \Lambda_{\partial B}}) = Z_{\Lambda_B}(\lambda_{\eta \cap (\Lambda_B)^c}) = Z_{\Lambda_B}(\lambda_{(\eta \cap \Lambda_H) \cup (\eta \cap \Lambda_{S \setminus \{\mathbf{v}\})}),$$

where the last equality comes from the fact that H and $S \setminus \{\mathbf{v}\}$ form a partitioning of $(B)^c$. Similarly, we obtain

$$Z_{\Lambda_{B \setminus \{\mathbf{v}\}}}(\lambda_{\eta \cap \Lambda_{\partial B \cup \{\mathbf{v}\}}}) = Z_{\Lambda_{B \setminus \{\mathbf{v}\}}}(\lambda_{\eta \cap (\Lambda_{B \setminus \{\mathbf{v}\})^c}) = Z_{\Lambda_{B \setminus \{\mathbf{v}\}}}(\lambda_{(\eta \cap \Lambda_H) \cup (\eta \cap \Lambda_S)}).$$

Thus, applying Lemma 6.1 with $\xi_1 = \gamma$ and $\xi_2 = \eta \cap \Lambda_H$ yields

$$\begin{aligned} & \min_{\gamma \subseteq (\delta_1 \mathbb{Z})^d \cap \Lambda_{H \cap \partial B}} \left\{ \frac{Z_{\Lambda_{B \setminus \{\mathbf{v}\}}}(\lambda_{\gamma \cup (\eta \cap \Lambda_S)})}{Z_{\Lambda_B}(\lambda_{\gamma \cup (\eta \cap \Lambda_{S \setminus \{\mathbf{v}\})})} \right\} \cdot \frac{Z_{\Lambda_B}(\lambda_{\eta \cap \Lambda_{\partial B}})}{Z_{\Lambda_{B \setminus \{\mathbf{v}\}}}(\lambda_{\eta \cap \Lambda_{\partial B \cup \{\mathbf{v}\}}})} \\ & \geq \exp\left(-a 3^{d_r d} e^{2br} (\lambda r^d + e^{\lambda 3^{d_r d}}) e^{-br \ell}\right). \end{aligned}$$

Moreover, following the same arguments as in Lemma 7.4, we have

$$e^{-e^{-\ell}/2} \cdot \min_{\gamma \subseteq (\delta_1 \mathbb{Z})^d \cap \Lambda_{H \cap \partial B}} \frac{Z_{\Lambda_{B \setminus \{\mathbf{v}\}}}(\lambda_{\gamma \cup (\eta \cap \Lambda_S)})}{Z_{\Lambda_B}(\lambda_{\gamma \cup (\eta \cap \Lambda_{S \setminus \{\mathbf{v}\})})} \leq \min_{\gamma \subseteq (\delta_1 \mathbb{Z})^d \cap \Lambda_{H \cap \partial B}} \frac{\hat{Z}(B \setminus \{\mathbf{v}\}, \gamma \cup (\eta \cap \Lambda_S), \delta_2)}{\hat{Z}(B, \gamma \cup (\eta \cap \Lambda_{S \setminus \{\mathbf{v}\})}, \delta_2)}.$$

Recalling the definition of $C_{e^{-\ell}}(S, \mathbf{v}, \eta)$ in Lemma 7.4 and noting that for $\ell \geq 2$ it holds that $b \cdot (\ell - 1)r \geq \frac{br}{2}\ell$, this implies

$$\begin{aligned} e^{-e^{-\ell}} C_{e^{-\ell}}(S, \mathbf{v}, \eta) \cdot \frac{Z_{\Lambda_B}(\lambda_{\eta \cap \Lambda_{\partial B}})}{Z_{\Lambda_B \setminus \{\mathbf{v}\}}(\lambda_{\eta \cap \Lambda_{\partial B \cup \{\mathbf{v}\}}})} &\geq \exp\left(-a3^{d_r d} e^{2br} (\lambda r^d + e^{\lambda 3^d r^d}) e^{-br\ell}\right) \cdot \exp\left(2e^{-\ell}\right) \\ &\geq 1 - a' e^{-b'\ell} \end{aligned}$$

for $b' = \min\{1, br\}$ and $a' = a3^{d_r d} e^{2br} (\lambda r^d + e^{\lambda 3^d r^d}) + 2$, which concludes the proof. \square

Lemma 7.9 allows us to control the success probability of the Bayes filter in terms of ℓ . This leads to the following statement.

Lemma 7.10. *Consider a hard-sphere model that exhibits (a, b) -strong spatial mixing up to λ . Suppose we run Algorithm 1 with $C(\cdot) = e^{-e^{-\ell}} C_{e^{-\ell}}(\cdot)$ as Bayes filter correction in line 5, and let $T = \inf\{t \in \mathbb{N}_0 \mid \mathcal{U}_t = \emptyset\}$ denote the number of iterations until the algorithm terminates. Then, for ℓ sufficiently large depending on a, b, r, λ and d , it holds that $\mathbb{E}[T] \leq 2|\mathcal{V}|$.*

Proof. We aim for applying Theorem B.2 to prove our claim. To this end, consider the process $(|\mathcal{U}_t|)_{t \in \mathbb{N}_0}$ and the filtration $(\mathcal{F}_t)_{t \in \mathbb{N}_0}$ defined by $\mathcal{F}_t = \sigma((X_j, \mathcal{U}_j, \mathbf{u}_j)_{j \leq t})$. Further, observe that our desired hitting time can equivalently be written as $T = \inf\{t \in \mathbb{N}_0 \mid |\mathcal{U}_t| \leq 0\}$. Since we are interested in the expectation of T , we only need to check assumptions a) and b) of Theorem B.2.

For a) of Theorem B.2, observe that $|\mathcal{U}_t| \geq 0$ for all $t \in \mathbb{N}_0$. For b), we prove that $\mathbb{E}[(|\mathcal{U}_t| - |\mathcal{U}_{t+1}|) \mathbb{1}_{T>t} \mid \mathcal{F}_t] \geq \frac{1}{2} \mathbb{1}_{T>t}$ if ℓ is sufficiently large. Since

$$\mathbb{E}[(|\mathcal{U}_t| - |\mathcal{U}_{t+1}|) \mathbb{1}_{T>t} \mid \mathcal{F}_t] = (|\mathcal{U}_t| - \mathbb{E}[|\mathcal{U}_{t+1}| \mid \mathcal{F}_t]) \mathbb{1}_{T>t}$$

it suffices to show that $\mathbb{E}[|\mathcal{U}_{t+1}| \mid \mathcal{F}_t] \mathbb{1}_{T>t} \leq (|\mathcal{U}_t| - \frac{1}{2}) \mathbb{1}_{T>t}$.

To simplify notation, we will omit the indicator of $T > t$ while still restricting ourselves to the setting where $\mathcal{U}_t \neq \emptyset$. Next, observe that, if $F_t = 1$, then $|\mathcal{U}_{t+1}| = |\mathcal{U}_t| - 1$. On the other hand, if $F_t = 0$, then

$$|\mathcal{U}_{t+1}| = |\mathcal{U}_t| + |\partial B| \leq |\mathcal{U}_t| + |B \cup \partial B| \leq |\mathcal{U}_t| + (2\ell + 3)^d,$$

where $B = B(\mathcal{U}_t, \mathbf{u}_t, \ell)$. Thus, it suffices if

$$(2\ell + 3)^d \cdot (1 - \mathbb{E}[F_t \mid \mathcal{F}_t]) - \mathbb{E}[F_t \mid \mathcal{F}_t] \leq -\frac{1}{2}.$$

Since, by Lemma 7.9,

$$\mathbb{E}[F_t \mid \mathcal{F}_t] \geq 1 - a' e^{-b'\ell},$$

for a' and b' only depending on a, b, λ, r and d , this is satisfied for ℓ sufficiently large, depending on a, b, λ, r and d . Applying Theorem B.2 then yields $\mathbb{E}[T] \leq 2|\mathcal{V}|$, which concludes the proof. \square

We conclude the following theorem.

Theorem 7.11. *Consider Algorithm 1 on a hard-sphere model with $C(\cdot) = e^{-e^{-\ell}} C_{e^{-\ell}}(\cdot)$ as Bayes filter correction in line 5. We can run the algorithm in almost-surely finite running time and, on termination, it outputs a sample from the hard-sphere Gibbs measure μ_λ on Λ . Moreover, if the hard-sphere model satisfies (a, b) -strong spatial mixing and if ℓ is chosen as a sufficiently large constant, depending on a, b, r, λ and d , then we can run the algorithm in expected time $O(|\Lambda|)$.*

Proof. For the first part of the statement, note that the correct output distribution follows directly from Theorem 5.2 and the fact that $e^{-e^{-\ell}} C_{e^{-\ell}}(\cdot)$ is a Bayes filter correction by Lemma 7.4. Let T denote that number of iterations of Algorithm 1, and let R_t denote the running time in iterations $t \in \mathbb{N}$. By Corollary 5.1 we know that T is almost surely finite. Moreover, it holds that $\mathbb{E}[R_t] = \mathbb{E}[\mathbb{E}[R_t \mid X_{t-1}, \mathcal{U}_{t-1}, \mathbf{u}_{t-1}]]$. Since by Corollary 7.8 $\mathbb{E}[R_t \mid X_{t-1}, \mathcal{U}_{t-1}, \mathbf{u}_{t-1}] \leq t(\ell, r, \lambda, d)$ for some function $t : \mathbb{R}_{\geq 0}^3 \times \mathbb{N}_0 \rightarrow \mathbb{R}_{\geq 0}$, it also holds that $\mathbb{E}[R_t] \leq t(\ell, r, \lambda, d)$. Consequently, R_t must be finite almost surely, and Algorithm 1 has almost surely finite running time.

For the second part of the statement, suppose the hard-sphere model satisfies (a, b) -strong spatial mixing up to λ . Observe that the expected running time of Algorithm 1 can be expressed as

$$\begin{aligned} \mathbb{E} \left[\sum_{t=1}^T R_t \right] &= \mathbb{E} \left[\sum_{t \geq 1} \mathbb{1}_{t \geq T} \mathbb{E}[R_t \mid X_{t-1}, \mathcal{U}_{t-1}, \mathbf{u}_{t-1}] \right] \\ &\leq t(\ell, r, \lambda, d) \mathbb{E} \left[\sum_{t \geq 1} \mathbb{1}_{t \geq T} \right] \\ &= t(\ell, r, \lambda, d) \mathbb{E}[T], \end{aligned}$$

where the first equality uses the fact that $\mathbb{1}_{t \geq T} = \mathbb{1}_{\mathcal{U}_{t-1} \neq \emptyset}$ is $\sigma(\mathcal{U}_{t-1})$ -measurable. By Lemma 7.10, we can choose ℓ sufficiently large, depending on a, b, r, λ and d only, such that $\mathbb{E}[T] \leq 2|\mathcal{V}| \in O(|\Lambda|)$, proving linear expected running time of the algorithm for any such choice of ℓ . \square

8 General repulsive potentials

We now extend our perfect sampling algorithm to the setting of more general bounded-range repulsive potentials ϕ . In contrast to the hard-sphere model, it is not clear how to perform the minimization task involved in constructing the Bayes filter in this setting. We will instead assume knowledge of the rate of strong spatial mixing for constructing the Bayes filter.

Lemma 8.1. *Let $S \subseteq \mathcal{V}$ be non-empty and $\mathbf{v} \in S$. Let $\eta \in \mathcal{N}_\Lambda$ be feasible and set $B = B(S, \mathbf{v}, \ell)$. Suppose $a, b > 0$ are such that ϕ satisfies (a, b) -strong spatial mixing up to λ , and set*

$$\delta = \delta(a, b) := \exp\left(-a3^d r^d e^{2br} \left(\lambda r^d + e^{\lambda 3^d r^d}\right) e^{-br\ell}\right)$$

and

$$C_{a,b}(S, \mathbf{v}, \eta) := \delta \cdot \frac{Z_{\Lambda_B \setminus \{\mathbf{v}\}}(\lambda_\eta \cap \Lambda_S)}{Z_{\Lambda_B}(\lambda_\eta \cap \Lambda_{S \setminus \{\mathbf{v}\}})}.$$

Then $C_{a,b}(S, \mathbf{v}, \eta)$ is a Bayes filter correction as in Definition 4.1. Moreover, it holds that

$$\exp\left(-2a3^d r^d e^{2br} \left(\lambda r^d + e^{\lambda 3^d r^d}\right) e^{-br\ell}\right) \leq C_{a,b}(S, \mathbf{v}, \eta) \cdot \frac{Z_{\Lambda_B}(\lambda_\eta \cap \Lambda_{\partial B})}{Z_{\Lambda_B \setminus \{\mathbf{v}\}}(\lambda_\eta \cap \Lambda_{\partial B \cup \{\mathbf{v}\}})} \leq 1.$$

Proof. We start by checking that $C_{a,b}(\cdot)$ is a Bayes filter correction. For the measurability, note that for all non-empty $S \subseteq \mathcal{V}$, $\mathbf{v} \in S$ and $\eta \in \mathcal{N}_\Lambda$ it holds that $C_{a,b}(S, \mathbf{v}, \eta)$ does only depend on $\eta \cap \Lambda_S$. Moreover, observe that

$$\begin{aligned} C_{a,b}(S, \mathbf{v}, \eta) &\geq \exp\left(-a3^d r^d e^{2br} \left(\lambda r^d + e^{\lambda 3^d r^d}\right) e^{-br\ell}\right) \cdot e^{-\lambda |\Lambda_{B \cup \partial B}|} \\ &\geq \exp\left(-a3^d r^d e^{2br} \left(\lambda r^d + e^{\lambda 3^d r^d}\right) e^{-br\ell}\right) \cdot \exp\left(-\lambda(2\ell + 3)^d r^d\right) \end{aligned}$$

uniformly in η . For the upper bound, we apply Lemma 6.1 with $\xi_1 = \emptyset$ and $\xi_2 = \xi$ to obtain

$$\frac{Z_{\Lambda_B \setminus \{\mathbf{v}\}}(\lambda_\eta \cap \Lambda_S)}{Z_{\Lambda_B}(\lambda_\eta \cap \Lambda_{S \setminus \{\mathbf{v}\}})} \cdot \inf_{\substack{\xi \in \mathcal{N}_{\Lambda_H} \\ \xi \cup (\eta \cap \Lambda_S) \text{ is feasible}}} \frac{Z_{\Lambda_B}(\lambda_{\xi \cup (\eta \cap \Lambda_{S \setminus \{\mathbf{v}\}})})}{Z_{\Lambda_B \setminus \{\mathbf{v}\}}(\lambda_{\xi \cup (\eta \cap \Lambda_S)})} \leq \exp\left(a3^d r^d e^{2br} \left(\lambda r^d + e^{\lambda 3^d r^d}\right) e^{-br\ell}\right).$$

Therefore, multiplying both sides with δ yields

$$C_{a,b}(S, \mathbf{v}, \eta) \leq \inf_{\substack{\xi \in \mathcal{N}_{\Lambda_H} \\ \xi \cup (\eta \cap \Lambda_S) \text{ is feasible}}} \frac{Z_{\Lambda_B \setminus \{\mathbf{v}\}}(\lambda_{\xi \cup (\eta \cap \Lambda_S)})}{Z_{\Lambda_B}(\lambda_{\xi \cup (\eta \cap \Lambda_{S \setminus \{\mathbf{v}\}})})}$$

as desired.

For the second part of the statement, note that by Lemma C.2 we have

$$Z_{\Lambda_B}(\lambda_\eta \cap \Lambda_{\partial B}) = Z_{\Lambda_B}(\lambda_\eta \cap (\Lambda_B)^c) = Z_{\Lambda_B}(\lambda_{(\eta \cap \Lambda_H) \cup (\eta \cap \Lambda_{S \setminus \{\mathbf{v}\}})}),$$

where the last equality comes from the fact that H and $S \setminus \{\mathbf{v}\}$ form a partitioning of $(B)^c$. Similarly, we obtain

$$Z_{\Lambda_B \setminus \{\mathbf{v}\}}(\lambda_\eta \cap \Lambda_{\partial B \cup \{\mathbf{v}\}}) = Z_{\Lambda_B \setminus \{\mathbf{v}\}}(\lambda_\eta \cap (\Lambda_B \setminus \{\mathbf{v}\})^c) = Z_{\Lambda_B \setminus \{\mathbf{v}\}}(\lambda_{(\eta \cap \Lambda_H) \cup (\eta \cap \Lambda_S)}).$$

Since $C_{a,b}(S, \mathbf{v}, \eta)$ is a Bayes filter correction, it follows that

$$C_{a,b}(S, \mathbf{v}, \eta) \cdot \frac{Z_{\Lambda_B}(\lambda_{\eta \cap \Lambda_{\partial B}})}{Z_{\Lambda_{B \setminus \{\mathbf{v}\}}}(\lambda_{\eta \cap \Lambda_{\partial B \cup \{\mathbf{v}\}}})} \leq 1.$$

Moreover, applying Lemma 6.1 with $\xi_1 = \emptyset$ and $\xi_2 = \eta \cap \Lambda_H$ yields

$$\frac{Z_{\Lambda_{B \setminus \{\mathbf{v}\}}}(\lambda_{\eta \cap \Lambda_S})}{Z_{\Lambda_B}(\lambda_{\eta \cap \Lambda_{S \setminus \{\mathbf{v}\}}})} \cdot \frac{Z_{\Lambda_B}(\lambda_{\eta \cap \Lambda_{\partial B}})}{Z_{\Lambda_{B \setminus \{\mathbf{v}\}}}(\lambda_{\eta \cap \Lambda_{\partial B \cup \{\mathbf{v}\}}})} \geq \exp\left(-a3^d r^d e^{2br} (\lambda r^d + e^{\lambda 3^d r^d}) e^{-br\ell}\right)$$

and multiplying both sides with δ proves the claim. \square

While the first part of Lemma 8.1 is sufficient to guarantee correctness of Algorithm 1, the second part of the statement will be useful for bounding the running by allowing us to control the probability of the event $F_t = 0$ for each iteration $t \in \mathbb{N}_0$. Similarly as in the setting of the hard-sphere model, we will not work directly with $C_{a,b}(S, \mathbf{v}, \eta)$, but rather use a slightly scaled version, which is a Bayes filter correction in its own right.

Corollary 8.2. *In the setting of Lemma 8.1, it holds that $e^{-e^{-\ell}} C_{a,b}(S, \mathbf{v}, r)$ is a Bayes filter correction, and there are constants a', b' , only depending on a, b, r, λ and d , such that*

$$1 - a'e^{-b'\ell} \leq e^{-e^{-\ell}} C_{a,b}(S, \mathbf{v}, \eta) \cdot \frac{Z_{\Lambda_B}(\lambda_{\eta \cap \Lambda_{\partial B}})}{Z_{\Lambda_{B \setminus \{\mathbf{v}\}}}(\lambda_{\eta \cap \Lambda_{\partial B \cup \{\mathbf{v}\}}})} \leq e^{-e^{-\ell}}.$$

Proof. Since $0 < e^{-e^{-\ell}} \leq 1$, it is obvious that $e^{-e^{-\ell}} C_{a,b}(S, \mathbf{v}, r)$ satisfies Definition 4.1. Now, note that for $\ell \geq 2$ it holds that $b(\ell - 1)r \geq \frac{br}{2}\ell$. The statement directly follows from Lemma 8.1 by setting $b' = \min\{1, br\}$ and $a' = 2a3^d r^d e^{2br} (\lambda r^d + e^{\lambda 3^d r^d}) + 1$, and observing that $\exp(-a'e^{-b'\ell}) \geq 1 - a'e^{-b'\ell}$. \square

Next, we focus on how to sample the Bayes filter, using $e^{-e^{-\ell}} C_{a,b}(\cdot)$ as in Corollary 8.2 as Bayes filter correction. In contrast to our approach for the hard-sphere model, we do not know how to compute $C_{a,b}(\cdot)$ directly. Again, we solve this problem using a Bernoulli factory.

Lemma 8.3. *Let $S \subseteq \mathcal{V}$ be non-empty, $\mathbf{v} \in S$ and $\eta \in \mathcal{N}_\Lambda$ be feasible, and set $B = B(S, \mathbf{v}, \ell)$. Suppose $a, b > 0$ are such that ϕ satisfies (a, b) -strong spatial mixing up to λ and let $C_{a,b}(\cdot)$ be as in Lemma 8.1. We can sample a Bernoulli random variable with success probability*

$$e^{-e^{-\ell}} C_{a,b}(S, \mathbf{v}, \eta) \cdot \frac{Z_{\Lambda_B}(\lambda_{\eta \cap \Lambda_{\partial B}})}{Z_{\Lambda_{B \setminus \{\mathbf{v}\}}}(\lambda_{\eta \cap \Lambda_{\partial B \cup \{\mathbf{v}\}}})}$$

with expected running time in $O(|\eta \cap \Lambda_{\{\mathbf{v}\} \cup \partial B}|)$, where the constants only depend on ℓ, r, λ and d .

Proof. Our goal is to use a Bernoulli factory of the form $\frac{p}{q}$ to perform this task. To bring the desired success probability into such a form, note that

$$\begin{aligned} Z_{\Lambda_B}(\lambda_{\eta \cap \Lambda_{\partial B}}) &= \frac{1}{\mu_{\lambda_{\eta \cap \Lambda_{\partial B}}} \mathbb{1}_{\Lambda_B}(\{\emptyset\})} \\ Z_{\Lambda_{B \setminus \{\mathbf{v}\}}}(\lambda_{\eta \cap \Lambda_{\partial B \cup \{\mathbf{v}\}}}) &= \frac{1}{\mu_{\lambda_{\eta \cap \Lambda_{\partial B \cup \{\mathbf{v}\}}}} \mathbb{1}_{\Lambda_{B \setminus \{\mathbf{v}\}}}(\{\emptyset\})}. \end{aligned}$$

Moreover, using Lemma C.2 we have

$$\begin{aligned} Z_{\Lambda_{B \setminus \{\mathbf{v}\}}}(\lambda_{\eta \cap \Lambda_S}) &= Z_{\Lambda_{B \setminus \{\mathbf{v}\}}}(\lambda_{\eta \cap \Lambda_{(\partial B \cap S) \cup \{\mathbf{v}\}}}) = \frac{1}{\mu_{\lambda_{\eta \cap \Lambda_{(\partial B \cap S) \cup \{\mathbf{v}\}}}} \mathbb{1}_{\Lambda_{B \setminus \{\mathbf{v}\}}}(\{\emptyset\})} \\ Z_{\Lambda_B}(\lambda_{\eta \cap \Lambda_{S \setminus \{\mathbf{v}\}}}) &= Z_{\Lambda_B}(\lambda_{\eta \cap \Lambda_{\partial B \cap S}}) = \frac{1}{\mu_{\lambda_{\eta \cap \Lambda_{\partial B \cap S}}} \mathbb{1}_{\Lambda_B}(\{\emptyset\})}. \end{aligned}$$

Finally, note that $e^{-e^{-\ell}} \in [0, 1]$ and, for $\delta = \delta$ as in Lemma 8.1, $\delta \in [0, 1]$. Thus, by setting

$$\begin{aligned} p &:= e^{-e^{-\ell}} \delta \cdot \mu_{\lambda_{\eta \cap \Lambda_{\partial B \cap S}}} \mathbb{1}_{\Lambda_B}(\{\emptyset\}) \cdot \mu_{\lambda_{\eta \cap \Lambda_{\partial B \cup \{\mathbf{v}\}}}} \mathbb{1}_{\Lambda_{B \setminus \{\mathbf{v}\}}}(\{\emptyset\}) \\ q &:= \mu_{\lambda_{\eta \cap \Lambda_{(\partial B \cap S) \cup \{\mathbf{v}\}}}} \mathbb{1}_{\Lambda_{B \setminus \{\mathbf{v}\}}}(\{\emptyset\}) \cdot \mu_{\lambda_{\eta \cap \Lambda_{\partial B}}} \mathbb{1}_{\Lambda_B}(\{\emptyset\}) \end{aligned}$$

we have $p \in [0, 1], q \in [0, 1]$ and

$$e^{-e^{-\ell}} C_{a,b}(S, \mathbf{v}, \eta) \cdot \frac{Z_{\Lambda_B}(\lambda_{\eta \cap \Lambda_{\partial B}})}{Z_{\Lambda_{B \setminus \{\mathbf{v}\}}}(\lambda_{\eta \cap \Lambda_{\partial B \cup \{\mathbf{v}\}}})} = \frac{p}{q}.$$

We are now going to use Lemma 3.1 to prove that we can obtain a sample from $\text{Ber}\left(\frac{p}{q}\right)$ within the desired expected running time. To this end, we need to provide a positive lower bound on $q - p$ and we need an efficient way for generating independent samples from $\text{Ber}(q)$ and $\text{Ber}(p)$.

For the lower bound, note that by Corollary 8.2 it holds that

$$e^{-e^{-\ell}} C_{a,b}(S, \mathbf{v}, \eta) \cdot \frac{Z_{\Lambda_B}(\lambda_{\eta \cap \Lambda_{\partial B}})}{Z_{\Lambda_{B \setminus \{\mathbf{v}\}}}(\lambda_{\eta \cap \Lambda_{\partial B \cup \{\mathbf{v}\}}})} \leq e^{-e^{-\ell}}.$$

Consequently, we have

$$\begin{aligned} q - p &\geq \left(1 - e^{-e^{-\ell}}\right) q \\ &= \left(1 - e^{-e^{-\ell}}\right) \cdot \left(Z_{\Lambda_{B \setminus \{\mathbf{v}\}}}(\lambda_{\eta \cap \Lambda_{(\partial B \cap S) \cup \{\mathbf{v}\}}}) \cdot Z_{\Lambda_B}(\lambda_{\eta \cap \Lambda_{\partial B}})\right)^{-1} \\ &\geq \left(1 - e^{-e^{-\ell}}\right) e^{-2\lambda|\Lambda_B|}. \end{aligned}$$

Using the upper bound $|\Lambda_B| \leq (2\ell + 1)^{d r^d}$ yields $q - p \geq (1 - e^{-e^{-\ell}}) e^{-\lambda(2\ell+1)^{d r^d}}$.

We proceed by arguing that we can obtain an oracle for $\text{Ber}(p)$ and $\text{Ber}(q)$ as required by Lemma 3.1. In particular, we focus on $\text{Ber}(p)$ since $\text{Ber}(q)$ can be treated analogously. Firstly, note that we can sample a Bernoulli random variable with success probability $e^{-e^{-\ell}} \delta$ in constant time, since we can compute it explicitly. It remains to argue that we can sample Bernoulli random variables with success probability $\mu_{\lambda_{\eta \cap \Lambda_{\partial B \cap S}} \mathbb{1}_{\Lambda_B}}(\{\emptyset\})$ and $\mu_{\lambda_{\eta \cap \Lambda_{\partial B \cup \{\mathbf{v}\}}} \mathbb{1}_{\Lambda_B \setminus \{\mathbf{v}\}}}(\{\emptyset\})$ in the desired running time. Again, we focus on $\mu_{\lambda_{\eta \cap \Lambda_{\partial B \cap S}} \mathbb{1}_{\Lambda_B}}(\{\emptyset\})$ and treat $\mu_{\lambda_{\eta \cap \Lambda_{\partial B \cup \{\mathbf{v}\}}} \mathbb{1}_{\Lambda_B \setminus \{\mathbf{v}\}}}(\{\emptyset\})$ analogously. To this end, note that for $Y \sim \mu_{\lambda_{\eta \cap \Lambda_{\partial B \cap S}} \mathbb{1}_{\Lambda_B}}$ it holds that $Y \cap (\Lambda_B)^c = \emptyset$ almost surely. Thus, it holds that

$$\mu_{\lambda_{\eta \cap \Lambda_{\partial B}} \mathbb{1}_{\Lambda_B}}(\{\emptyset\}) = \mu_{\lambda_{\eta \cap \Lambda_{\partial B}} \mathbb{1}_{\Lambda_B}}[\Lambda_B](\{\emptyset\}),$$

and, for $Y \sim \mu_{\lambda_{\eta \cap \Lambda_{\partial B}} \mathbb{1}_{\Lambda_B}}[\Lambda_B]$, it holds that $\mathbb{1}_{Y=\emptyset}$ is a Bernoulli random variable with success probability $\mu_{\lambda_{\eta \cap \Lambda_{\partial B}} \mathbb{1}_{\Lambda_B}}(\{\emptyset\})$. By Lemma 7.5, we can obtain independent samples from $\mu_{\lambda_{\eta \cap \Lambda_{\partial B}} \mathbb{1}_{\Lambda_B}}[\Lambda_B]$, each in expected time at most $(\lambda|\Lambda_B| + |\eta \cap \Lambda_{\partial B}|) \cdot \lambda|\Lambda_B| e^{\lambda|\Lambda_B|}$. Noting that $|\Lambda_B| \leq (2\ell + 1)^{d r^d}$ and that $|\eta \cap \Lambda_{\partial B}| \leq |\eta \cap \Lambda_{\{\mathbf{v}\} \cup \partial B}|$ yields an expected running time of $O(|\eta \cap \Lambda_{\{\mathbf{v}\} \cup \partial B}|)$. Applying the same argument to sample a Bernoulli random variable with success probability $\mu_{\lambda_{\eta \cap \Lambda_{\partial B \cup \{\mathbf{v}\}}} \mathbb{1}_{\Lambda_B \setminus \{\mathbf{v}\}}}(\{\emptyset\})$ yields an oracle for $\text{Ber}(p)$ with expected running time in $O(|\eta \cap \Lambda_{\{\mathbf{v}\} \cup \partial B}|)$. Finally, applying the same procedure for $\text{Ber}(q)$ and using Lemma 3.1 concludes the proof. \square

We obtain the following bound for the running time of each iteration.

Corollary 8.4. *Let $a, b > 0$ be such that ϕ satisfies (a, b) -strong spatial mixing up to λ . Suppose we run Algorithm 1 on ϕ with $C(\cdot) = e^{-e^{-\ell}} C_{a,b}(\cdot)$ as Bayes filter correction in line 5, and let R_t denote the running time for iteration $t \in \mathbb{N}$. Then, for all $t \in \mathbb{N}$, it holds that*

$$\mathbb{E}[R_t \mid X_{t-1}, \mathcal{U}_{t-1}, \mathbf{u}_{t-1}] \leq O(|X_{t-1} \cap \Lambda_{\{\mathbf{u}_{t-1}\} \cup \partial B}|),$$

where $B = B(\mathcal{U}_{t-1}, \mathbf{u}_{t-1}, \ell)$, and the constants in the asymptotic notation only depend on ℓ , r , λ and d .

Proof. Set $B = B(\mathcal{U}_{t-1}, \mathbf{u}_{t-1}, \ell)$, and note that the bulk of the running time in each iteration of Algorithm 1 is due to sampling the Bayes filter in line 5 and updating the point configuration on Λ_B in line 7.

For line 5, note that X_{t-1} is almost surely feasible by Lemma 4.3. Thus, Lemma 8.3 yields that the expected time for sampling the Bayes filter, conditioned on $X_{t-1}, \mathcal{U}_{t-1}$ and \mathbf{u}_{t-1} , almost surely bounded by some function in $O(|X_{t-1} \cap \Lambda_{\{\mathbf{u}_{t-1}\} \cup \partial B}|)$. Further, for line 7, we can use Lemma 7.5 to bound the expected time for sampling from $\mu_{\lambda_{X_{t-1} \cap (\Lambda_B)^c} \mathbb{1}_{\Lambda_B}}[\Lambda_B]$ is bounded by $(\lambda|\Lambda_B| + |X_{t-1} \cap \Lambda_{\partial B}|) \cdot \lambda|\Lambda_B| e^{\lambda|\Lambda_B|} \in O(|X_{t-1} \cap \Lambda_{\{\mathbf{u}_{t-1}\} \cup \partial B}|)$ as desired. \square

Next, we derive a bound on the expected number of iterations of Algorithm 1 given strong spatial mixing.

Lemma 8.5. *Let $a, b > 0$ be such that ϕ satisfies (a, b) -strong spatial mixing up to λ . Suppose we run Algorithm 1 with $C(\cdot) = e^{-e^{-\ell}} C_{a,b}(\cdot)$ as Bayes filter correction in line 5, and let $T = \inf\{t \in \mathbb{N}_0 \mid \mathcal{U}_t = \emptyset\}$ denote the number of iterations until the algorithm terminates. Then, for ℓ sufficiently large depending on a, b, r, λ and d , it holds that $\mathbb{E}[T] \leq 2|\mathcal{V}|$, and for all $k \geq 4|\mathcal{V}|$ it holds that $\mathbb{P}[T \geq k] \leq \exp(-\frac{k}{\alpha})$ for some constant $\alpha \in \mathbb{R}_{>0}$ that only depends on ℓ, d and r .*

Proof. We aim for applying Theorem B.2 to prove our claim. For bounding $\mathbb{E}[T]$, we proceed analogously as in the proof of Lemma 7.10. In particular, we consider the process $(|\mathcal{U}_t|)_{t \in \mathbb{N}_0}$ with the filtration $(\mathcal{F}_t)_{t \in \mathbb{N}_0}$ defined by $\mathcal{F}_t = \sigma((X_j, \mathcal{U}_j, \mathbf{u}_j)_{j \leq t})$ and rewrite our desired hitting time as $T = \inf\{t \in \mathbb{N}_0 \mid |\mathcal{U}_t| \leq 0\}$. Since a) of Theorem B.2 is trivially satisfied, we only need to check b). Using the lower bound from Corollary 8.2 and the same arguments as in the proof of Lemma 7.10, we can show that $\mathbb{E}[(|\mathcal{U}_t| - |\mathcal{U}_{t+1}|) \mathbf{1}_{T > t} \mid \mathcal{F}_t] \geq \frac{1}{2} \mathbf{1}_{T > t}$ if ℓ is sufficiently large, depending on a, b, r, λ and d . Thus, applying the first part of Theorem B.2 proves our bound on $\mathbb{E}[T]$.

To obtain the tail bound on T , we apply the second part of Theorem B.2. For c), note that $|\mathcal{U}_0| = |\mathcal{V}|$ and, for d), observe that

$$\| |\mathcal{U}_t| - |\mathcal{U}_{t-1}| \| \leq \max\left\{1, (2\ell + 3)^d r^{2d}\right\}.$$

Thus, setting $\alpha \geq 64 \cdot \max\{1, (2\ell + 3)^{2d} r^{2d}\}$ concludes the proof. \square

Note that, in contrast to Corollary 7.8, Corollary 8.4 does not give a deterministic bound on the running time of each iteration. Due to potential dependencies between the running time of each iteration and termination of the algorithm, it is unclear if we can simply apply Wald's equation to derive the total running time of the algorithm. Instead, we will use a more subtle argument for this. As a first ingredient, we need upper-bound the probability of ever observing a large number of points in any box $\Lambda_{\mathbf{v}}$ for $\mathbf{v} \in \mathcal{V}$ up to a given iteration $k \in \mathbb{N}_0$.

Lemma 8.6. *There is a constant $\alpha \in \mathbb{R}_{\geq 1}$, only depending on ℓ, d and r , such that, for all $k \in \mathbb{N}_0$, all $\gamma > 1$ and all $x \geq e^\gamma \lambda r^d$, it holds that*

$$\mathbb{P}\left[\bigcup_{t=0}^k \bigcup_{\mathbf{v} \in \mathcal{V}} \{|X_t \cap \Lambda_{\mathbf{v}}| \geq x\} \right] \leq (k+1)\alpha e^{-(\gamma-1)x}.$$

Proof. Let $\alpha \geq (2\ell + 1)^d$. We show the statement via induction over k . First, for $k = 0$, note that $X_0 = \emptyset$. Thus, we have that the left-hand side is 0 for all $x > 0$ and the right-hand side is at least 1 for $x = 0$, proving the base case.

Next, assume the statement holds for some fixed $k \in \mathbb{N}_{\geq 0}$. Let $\hat{B} \subseteq \mathcal{V}$ be the set of boxes updated in iteration $k + 1$ (possibly the empty set). Formally, that is $\hat{B} =$

$\{\mathbf{v} \in \mathcal{V} \mid X_{k+1} \cap \Lambda_{\mathbf{v}} \neq X_k \cap \Lambda_{\mathbf{v}}\}$. By the induction hypothesis and union bound, it holds that

$$\begin{aligned} \mathbb{P} \left[\bigcup_{t=0}^{k+1} \bigcup_{\mathbf{v} \in \mathcal{V}} \{|X_t \cap \Lambda_{\mathbf{v}}| \geq x\} \right] &= \mathbb{P} \left[\bigcup_{t=0}^k \bigcup_{\mathbf{v} \in \mathcal{V}} \{|X_t \cap \Lambda_{\mathbf{v}}| \geq x\} \cup \left\{ \exists \mathbf{v} \in \hat{B} \text{ s.t. } |X_{k+1} \cap \Lambda_{\mathbf{v}}| \geq x \right\} \right] \\ &\leq \mathbb{P} \left[\bigcup_{t=0}^k \bigcup_{\mathbf{v} \in \mathcal{V}} \{|X_t \cap \Lambda_{\mathbf{v}}| \geq x\} \right] + \mathbb{P} \left[\left\{ \exists \mathbf{v} \in \hat{B} \text{ s.t. } |X_{k+1} \cap \Lambda_{\mathbf{v}}| \geq x \right\} \right] \\ &\leq (k+1)\alpha \cdot e^{-(\gamma-1)x} + \mathbb{P} \left[\left\{ \exists \mathbf{v} \in \hat{B} \text{ s.t. } |X_{k+1} \cap \Lambda_{\mathbf{v}}| \geq x \right\} \right]. \end{aligned}$$

It now suffices to show that

$$\mathbb{P} \left[\left\{ \exists \mathbf{v} \in \hat{B} \text{ s.t. } |X_{k+1} \cap \Lambda_{\mathbf{v}}| \geq x \right\} \right] \leq \alpha \cdot \mathbb{P}[Y \geq x].$$

First, note that $\hat{B} = \emptyset$ if $F_k = 0$. Thus, we have

$$\mathbb{P} \left[\left\{ \exists \mathbf{v} \in \hat{B} \text{ s.t. } |X_{k+1} \cap \Lambda_{\mathbf{v}}| \geq x \right\} \mid F_k = 0 \right] \leq \alpha \cdot \mathbb{P}[Y \geq x]$$

for all $x \in \mathbb{R}_{\geq 0}$. Now, fix $S \subseteq \mathcal{V}$ and $\mathbf{w} \in S$ such that $\mathbb{P}[\mathcal{U}_k = S, \mathbf{u}_k = \mathbf{w}, F_k = 1] > 0$. Set $B = B(S, \mathbf{w}, \ell)$ and observe that, given $\mathcal{U}_k = S, \mathbf{u}_k = \mathbf{u}$ and $F_k = 1$, it holds that $\hat{B} = B$. Using union bound, we have

$$\begin{aligned} &\mathbb{P} \left[\left\{ \exists \mathbf{v} \in \hat{B} \text{ s.t. } |X_{k+1} \cap \Lambda_{\mathbf{v}}| \geq x \right\} \mid \mathcal{U}_k = S, \mathbf{u}_k = \mathbf{w}, F_k = 1 \right] \\ &= \mathbb{P} \left[\bigcup_{\mathbf{v} \in B} \{|X_{k+1} \cap \Lambda_{\mathbf{v}}| \geq x\} \mid \mathcal{U}_k = S, \mathbf{u}_k = \mathbf{w}, F_k = 1 \right] \\ &\leq \sum_{\mathbf{v} \in B} \mathbb{P}[|X_{k+1} \cap \Lambda_{\mathbf{v}}| \geq x \mid \mathcal{U}_k = S, \mathbf{u}_k = \mathbf{w}, F_k = 1]. \end{aligned}$$

Given $\mathcal{U}_k = S, \mathbf{u}_k = \mathbf{w}$ and $F_k = 1$, it holds that $X_{k+1} \cap \Lambda_B$ is sampled from $\mu_{\lambda_{\eta \cap (\Lambda_B)^c} \mathbf{1}_{\Lambda_B}}[\Lambda_B]$ for some feasible $\eta \in \mathcal{N}_{\Lambda}$. Therefore, for all $\mathbf{v} \in B$, $|X_{k+1} \cap \Lambda_{\mathbf{v}}|$ is dominated by a Poisson random variable Y with parameter $\lambda|\Lambda_{\mathbf{v}}| \leq \lambda r^d$. Further, observing that $|B| \leq \alpha$ yields

$$\mathbb{P} \left[\left\{ \exists \mathbf{v} \in \hat{B} \text{ s.t. } |X_{k+1} \cap \Lambda_{\mathbf{v}}| \geq x \right\} \mid \mathcal{U}_k = S, \mathbf{u}_k = \mathbf{w}, F_k = 1 \right] \leq \alpha \cdot \mathbb{P}[Y \geq x].$$

Using the law of total expectation and Corollary B.4, we obtain

$$\mathbb{P} \left[\left\{ \exists \mathbf{v} \in \hat{B} \text{ s.t. } |X_{k+1} \cap \Lambda_{\mathbf{v}}| \geq x \right\} \right] \leq \alpha \cdot e^{-(\gamma-1)x},$$

which proves the claim. \square

Using Lemma 8.6, we derive our main result on perfect sampling for repulsive bounded-range potentials based on Algorithm 1.

Theorem 8.7. *Let $a, b > 0$ be such that ϕ satisfies (a, b) -strong spatial mixing up to λ . Consider Algorithm 1 with $C(\cdot) = e^{-e^{-\ell}} C_{a,b}(\cdot)$ as Bayes filter correction in line 5. On termination, the algorithm outputs a sample from the Gibbs measure μ_λ on Λ . Moreover, if ℓ is chosen as a sufficiently large constant, depending on a, b, r, λ and d , then we can run the algorithm in expected time $\tilde{O}(|\Lambda|)$.*

Proof. If ϕ satisfies (a, b) -strong spatial mixing up to λ , we know by Corollary 8.2 that $e^{-e^{-\ell}} C_{a,b}(\cdot)$ is a Bayes filter correction. Thus, the first part of the statement follows from Theorem 5.2.

Next, let ℓ be chosen as a sufficiently large constant, depending on a, b, r, λ and d to satisfy Lemma 8.5. We rewrite the running time of Algorithm 1 as

$$\begin{aligned} \mathbb{E} \left[\sum_{t=1}^T R_t \right] &= \mathbb{E} \left[\mathbb{E} \left[\sum_{t=1}^{\infty} \mathbb{1}_{T \geq t} R_t \mid \mathcal{U}_{t-1}, X_{t-1}, \mathbf{u}_{t-1} \right] \right] \\ &= \mathbb{E} \left[\sum_{t=1}^{\infty} \mathbb{1}_{T \geq t} \mathbb{E}[R_t \mid \mathcal{U}_{t-1}, X_{t-1}, \mathbf{u}_{t-1}] \right] \end{aligned}$$

Moreover, by Corollary 8.4, there are constants a_1, a_2 , only depending on ℓ, d, λ and r , such that

$$\mathbb{E}[R_t \mid \mathcal{U}_{t-1}, X_{t-1}, \mathbf{u}_{t-1}] \leq \alpha_1 |X_{t-1} \cap \Lambda_{\{\mathbf{u}_{t-1}\} \cup \partial B_t}| + \alpha_2$$

where $B_t = B(\mathcal{U}_{t-1}, \mathbf{u}_{t-1}, \ell)$. Thus, we have

$$\mathbb{E} \left[\sum_{t=1}^T R_t \right] \leq \alpha_1 \mathbb{E} \left[\sum_{t=1}^{\infty} \mathbb{1}_{T \geq t} |X_{t-1} \cap \Lambda_{\{\mathbf{u}_{t-1}\} \cup \partial B_t}| \right] + \alpha_2 \mathbb{E} \left[\sum_{t=1}^{\infty} \mathbb{1}_{T \geq t} \right].$$

Note that by Lemma 8.5

$$\sum_{t=1}^{\infty} \mathbb{E}[\mathbb{1}_{T \geq t}] = \sum_{t'=0}^{\infty} \mathbb{P}[\mathbb{1}_{T > t'}] = \mathbb{E}[T] \in O(|\Lambda|).$$

It remains to bound

$$\mathbb{E} \left[\sum_{t=1}^{\infty} \mathbb{1}_{T \geq t} \cdot |X_{t-1} \cap \Lambda_{\{\mathbf{u}_{t-1}\} \cup \partial B_t}| \right].$$

To this end, write $W = \sum_{t=1}^{\infty} \mathbb{1}_{T \geq t} \cdot |X_{t-1} \cap \Lambda_{\{\mathbf{u}_{t-1}\} \cup \partial B_t}|$. Since W is non-negative, we have

$$\mathbb{E}[W] = \int_{\mathbb{R}_{\geq 0}} \mathbb{P}[W > w] dw \leq \hat{w} + \int_{\hat{w}}^{\infty} \mathbb{P}[W > w] dw$$

for every $\hat{w} \in \mathbb{R}_{\geq 0}$. Next, observe that for every $t \in \mathbb{N}$

$$|\{\mathbf{u}_{t-1}\} \cup \partial B_t| \leq |B_t \cup \partial B_t| \leq (2\ell + 1)^d =: \alpha_3.$$

Thus, for every $w \in \mathbb{R}_{\geq 0}$ and every $k : \mathbb{R}_{\geq 0} \rightarrow \mathbb{N}$ the following holds: if $W > w$, then $T > k(w)$ or there is a time point $t \leq k(w)$ and a box $\mathbf{v} \in \mathcal{V}$ such that $|X_{t-1} \cap \Lambda_{\mathbf{v}}| > \frac{w}{k(w)\alpha_3}$. Consequently, we have

$$\begin{aligned} \mathbb{P}[W > w] &\leq \mathbb{P}\left[\{T > k(w)\} \cup \bigcup_{t=1}^{k(w)} \bigcup_{\mathbf{v} \in \mathcal{V}} \left\{|X_{t-1} \cap \Lambda_{\mathbf{v}}| > \frac{w}{k(w)\alpha_3}\right\}\right] \\ &\leq \mathbb{P}[T > k(w)] + \mathbb{P}\left[\bigcup_{t=1}^{k(w)} \bigcup_{\mathbf{v} \in \mathcal{V}} \left\{|X_{t-1} \cap \Lambda_{\mathbf{v}}| > \frac{w}{k(w)\alpha_3}\right\}\right]. \end{aligned}$$

Moreover, applying Lemma 8.6 shows that there is a constant α_4 , only depending on ℓ , r and d , such that for all $w \geq e^2\alpha_3\lambda r^d k(w)$ it holds that

$$\mathbb{P}[W > w] \leq \mathbb{P}[T \geq k(w)] + (k(w) - 1)\alpha_4 \exp\left(-\frac{w}{k(w)\alpha_3}\right),$$

Now, suppose that k is measurable, we get for $\hat{w} \in \mathbb{R}_{\geq 0}$

$$\mathbb{E}[W] \leq \hat{w} + \int_{\hat{w}}^{\infty} \mathbb{P}[T \geq k(w)] dw + \alpha_4 \int_{\hat{w}}^{\infty} k(w) \exp\left(-\frac{w}{k(w)\alpha_3}\right) dw.$$

We now claim that for a suitable choice of \hat{w} and k the right-hand side is in $\tilde{O}(|\Lambda|)$.

To this end, let $k(w) = \max\{\lceil 4e|\mathcal{V}| \rceil, \lceil 2\alpha_5 \ln(w) \rceil\}$, where α_5 is the constant from the tail bound in Lemma 8.5, which only depends on ℓ , r and d . Moreover, choose $\hat{w} \geq 1$ sufficiently large such that for all $w \geq \hat{w}$ it holds that $w \geq k(w)$, $w \geq e^2\alpha_3\lambda r^d k(w)$ and $\frac{w}{\ln(w)} \geq 3\alpha_3 k(w)$.

Note that this can be achieved for some $\hat{w} \in \tilde{O}(|\mathcal{V}|)$.

For our choice of $k(w)$, Lemma 8.5 yields

$$\int_{\hat{w}}^{\infty} \mathbb{P}[T \geq k(w)] dw \leq \int_{\hat{w}}^{\infty} w^{-2} dw,$$

which is bounded by 1 for $\hat{w} \geq 1$. Moreover, for $w \geq k(w)$ and $\frac{w}{\ln(w)} \geq 3\alpha_3 k(w)$ we have

$$\exp\left(-\frac{w}{k(w)\alpha_3}\right) \leq \frac{1}{k(w)w^2}$$

and thus

$$\int_{\hat{w}}^{\infty} k(w) \exp\left(-\frac{w}{k(w)\alpha_3}\right) dw \leq \int_{\hat{w}}^{\infty} w^{-2} dw \leq 1.$$

Consequently, we have

$$\mathbb{E}[W] \leq \hat{w} + O(1) \in \tilde{O}(|\mathcal{V}|)$$

and, since $|\mathcal{V}| \in O(|\Lambda|)$, this concludes the proof. \square

Remark 8.8. *Since Theorem 8.7 requires knowing constants a, b such that the Point process satisfies (a, b) -strong spatial mixing (in contrast to Theorem 7.11), we can use these constant to compute a sufficiently large value for ℓ . Elementary calculations suggest to choose $\ell > \max\left\{\frac{1+7r}{2}, 16a'b'^2\right\}$, where $b' = \min\{1, br\}$ and $a' = 2a3^d r^d e^{2br} (\lambda r^d + e^{\lambda 3^d r^d}) + 1$.*

9 Bernoulli Factories

In this section we prove Lemma 3.1, showing how to sample a random variable from $\text{Ber}\left(\frac{p}{q}\right)$ given access to a $\text{Ber}(p)$ and $\text{Ber}(q)$ sampler, when $q - p > \epsilon$. This happens in the following three steps.

The first step is to sample a random variable according to $\text{Ber}\left(\frac{1-(q-p)}{2}\right)$ given access to $\text{Ber}(p)$ and $\text{Ber}(q)$. This is achieved by Algorithm 3 bellow. It is easy to verify that this

Algorithm 3: $\text{Ber}\left(\frac{1-(q-p)}{2}\right)$ from $\text{Ber}(p)$ and $\text{Ber}(q)$

```

1 Draw  $u \sim \text{Ber}(1/2)$ 
2 if  $u = 1$  then
3   Draw  $y \sim \text{Ber}(q)$ 
4   return  $1 - y$ 
5 else
6   Draw  $x \sim \text{Ber}(p)$ 
7   return  $x$ 

```

algorithm returns 1 with the correct probability.

For the second step, let $\varrho = \frac{1-(q-p)}{2}$. From Algorithm 3 we now assume to have access to a $\text{Ber}(\varrho)$ random variable. The next step is to use Huber's algorithm [37] and obtain a $\text{Ber}(2\varrho)$. For the algorithm to work within the required run-time guarantees, we need $\varrho < \frac{1-\epsilon}{2}$, which holds since we assumed $q - p > \epsilon$. For convenience, we provide the pseudocode of Huber's algorithm in Algorithm 4. The correctness of the algorithm can be found in [37, Section 2.3].

Algorithm 4 now gives us access to a $\text{Ber}(2\varrho)$ sampler and consequently to a $\text{Ber}(1 - 2\varrho) = \text{Ber}(q - p)$ sampler, simply by flipping the returned bit. The final step is to sample from $\text{Ber}\left(\frac{p}{q}\right)$ when given access to a sampler for $\text{Ber}(q - p)$. This is done via Algorithm 5.

Regarding the correctness of Algorithm 5 note that, within a single while-loop, the probability the algorithm returns 1 is $p/2$, while the probability that the algorithm enters the while-loop again is $1 - q/2$. Conditioned on the fact that the algorithm will terminate, we observe that Algorithm 5 returns 1 with probability p/q .

We are now ready to prove Lemma 3.1, whose statement we repeat here for convinience.

Algorithm 4: Ber(2ϱ) from Ber(ϱ)

```
1  $\epsilon \leftarrow \min\{\epsilon, 0.644\}$ ,  $k \leftarrow \frac{23}{5\epsilon}$ ,  $i \leftarrow 1$ ,  $R \leftarrow 1$ ,  $C \leftarrow 2$ 
2 while  $i \neq 0$  and  $R \neq 0$  do
3   while  $0 < i < k$  do
4     Draw  $u \sim \text{Ber}(\varrho)$ 
5     Draw  $g \sim \text{Geo}(\frac{C-1}{C})$ 
6      $i \leftarrow i - 1 + (1 - u)g$ 
7   if  $i \geq k$  then
8     Draw  $R \sim \text{Ber}\left(\left(1 + \frac{\epsilon}{2}\right)^{-i}\right)$ 
9     if  $R = 0$  then
10      return 0
11     $C \leftarrow C(1 + \frac{\epsilon}{2})$ ,  $\epsilon \leftarrow \frac{\epsilon}{2}$ ,  $k \leftarrow 2k$ 
12 return 1
```

Algorithm 5: Ber($\frac{p}{q}$) from Ber(p) and Ber($q - p$)

```
1 Set  $b = -1$ 
2 while  $b = -1$  do
3   Draw  $u_1 \sim \text{Ber}(1/2)$ 
4   if  $u_1 = 1$  then
5     Draw  $x \sim \text{Ber}(p)$ 
6     if  $x = 1$  then
7       Set  $b = 1$ 
8   else
9     Draw  $y \sim \text{Ber}(q - p)$ 
10    if  $y = 1$  then
11      Set  $b = 0$ 
12 return  $b$ 
```

Lemma 3.1. Fix some $p, q \in [0, 1]$ such that $q - p \geq \epsilon$ for some $\epsilon > 0$. Further assume that we have oracle access to a sampler from $\text{Ber}(p)$ and $\text{Ber}(q)$ in the following sense:

1. every sample from $\text{Ber}(p)$ (resp. $\text{Ber}(q)$) is independent from all previous samples;
2. the expected running time for obtaining a sample from $\text{Ber}(p)$ (resp. $\text{Ber}(q)$), conditioned on previously obtained samples, is uniformly bounded by some $t \in \mathbb{R}_{\geq 0}$.

Then we can sample from $\text{Ber}\left(\frac{p}{q}\right)$ in $O(t\epsilon^{-2})$ expected time.

Proof. We use Algorithm 5 which calls Algorithm 4, which in turn calls Algorithm 3, as we explained above. For simplicity, we may assume that, every time we sample $\text{Ber}(p)$, we also sample from $\text{Ber}(q)$ and vice versa. Thus, let $(X_i)_{i \in \mathbb{N}} \in (\{0, 1\}^2)^{\mathbb{N}}$ be a sequence of samples from the product distribution $(\text{Ber}(p) \otimes \text{Ber}(q))^{\otimes \mathbb{N}}$ and assume that $X_i = (X_i^{(p)}, X_i^{(q)})$ is the outcome of the i th time the algorithm samples from $\text{Ber}(p)$ and $\text{Ber}(q)$. Moreover, let $(T_i)_{i \in \mathbb{N}}$ be the running time for obtaining the X_i using the assumed oracle, and let N denote the total number of samples from $\text{Ber}(p)$ and $\text{Ber}(q)$ that Algorithm 5 requires.

Our goal is to show that $\mathbb{E}\left[\sum_{i=1}^N T_i\right] \in O(t\epsilon^{-2})$. To this end, note that, for every $i \in \mathbb{N}$, the event $N \geq i$ does only depend on the sequence $(X_i)_{i \in \mathbb{N}}$ via the subsequence $(X_j)_{j < i}$. Moreover, by our assumptions on the oracle, it holds that $\mathbb{E}[T_i \mid (X_j)_{j < i}] \leq 2t$, where the factor of 2 comes from the fact that we sample both $\text{Ber}(p)$ and $\text{Ber}(q)$. Thus, by Wald's equation, we have

$$\mathbb{E}\left[\sum_{i=1}^N T_i\right] = \mathbb{E}\left[\sum_{i \geq 1} \mathbf{1}_{N \geq i} \mathbb{E}[T_i \mid (X_j)_{j < i}]\right] \leq 2t\mathbb{E}[N].$$

It remains to show that $\mathbb{E}[N] \in O(\epsilon^{-2})$.

Algorithm 5 will do $2/q \in O(\epsilon^{-1})$ while-loops in expectation, as each while-loop terminates with probability $q/2$. Furthermore, each while-loop calls either $\text{Ber}(p)$ or Algorithm 4, both with probability $1/2$. As the total number of loops is determined by the outcome of the final loop, we can use Wald's equation again to get that $\mathbb{E}[N] \in O(\epsilon^{-1} + \epsilon^{-1}\mathbb{E}[N'])$, where N' is the number of X_i samples that Algorithm 4 requires. Observe (from the pseudocode of Algorithm 3) that each $\text{Ber}(\rho)$ -call Algorithm 4 requires only a single X_i sample. From [37, Theorem 1.1] we get that Algorithm 4 requires at most $O(\epsilon^{-1})$ $\text{Ber}(\rho)$ samples in expectation, which implies that $\mathbb{E}[N'] \in O(\epsilon^{-1})$. This concludes the proof of the lemma. \square

Acknowledgments

We thank Mark Jerrum for very helpful discussions on this topic. Konrad Anand was funded by a studentship from Queen Mary, University of London. Andreas Göbel was funded by the project PAGES (project No. 467516565) of the German Research Foundation (DFG). Marcus Pappik was funded by the HPI Research School on Data Science and Engineering. Will Perkins was supported in part by NSF grant DMS-2309958.

References

- [1] M. Aizenman and R. Holley. Rapid convergence to equilibrium of stochastic Ising models in the Dobrushin Shlosman regime. *Percolation theory and ergodic theory of infinite particle systems*, pages 1–11, 1987.
- [2] B. J. Alder and T. E. Wainwright. Phase transition for a hard sphere system. *The Journal of Chemical Physics*, 27(5):1208–1209, 1957.
- [3] K. Anand and M. Jerrum. Perfect sampling in infinite spin systems via strong spatial mixing. *SIAM Journal on Computing*, 51(4):1280–1295, 2022.
- [4] S. Asmussen, P. W. Glynn, and H. Thorisson. Stationarity detection in the initial transient problem. *ACM Transactions on Modeling and Computer Simulation (TOMACS)*, 2(2):130–157, 1992.
- [5] E. P. Bernard and W. Krauth. Two-step melting in two dimensions: first-order liquid-hexatic transition. *Physical Review Letters*, 107(15):155704, 2011.
- [6] E. P. Bernard, W. Krauth, and D. B. Wilson. Event-chain Monte Carlo algorithms for hard-sphere systems. *Physical Review E*, 80(5):056704, 2009.
- [7] S. Betsch and G. Last. On the uniqueness of Gibbs distributions with a non-negative and subcritical pair potential. *arXiv preprint arXiv:2108.06303*, 2021.
- [8] S. Bhandari and S. Chakraborty. Improved bounds for perfect sampling of k-colorings in graphs. In *Proceedings of the 52nd Annual ACM SIGACT Symposium on Theory of Computing*, pages 631–642, 2020.
- [9] Z. Chen, K. Liu, N. Mani, and A. Moitra. Strong spatial mixing for colorings on trees and its algorithmic applications. *arXiv preprint arXiv:2304.01954*, 2023.
- [10] H. T. Christoph. Disagreement percolation for the hard-sphere model. *Electronic Journal of Probability*, 24:1–22, 2019.
- [11] E. Çinlar. *Probability and stochastics*, volume 261. Springer, 2011.
- [12] D. Dereudre. Introduction to the theory of Gibbs point processes. In *Stochastic Geometry*, pages 181–229. Springer, 2019.
- [13] P. Diaconis. The Markov Chain Monte Carlo revolution. *Bulletin of the American Mathematical Society*, 46(2):179–205, 2009.
- [14] S. Dughmi, J. Hartline, R. D. Kleinberg, and R. Niazadeh. Bernoulli factories and black-box reductions in mechanism design. *Journal of the ACM (JACM)*, 68(2):1–30, 2021.

- [15] M. Dyer, A. Sinclair, E. Vigoda, and D. Weitz. Mixing in time and space for lattice spin systems: A combinatorial view. *Random Structures & Algorithms*, 24(4):461–479, 2004.
- [16] M. Engel, J. A. Anderson, S. C. Glotzer, M. Isobe, E. P. Bernard, and W. Krauth. Hard-disk equation of state: First-order liquid-hexatic transition in two dimensions with three simulation methods. *Physical Review E*, 87(4):042134, 2013.
- [17] S. Felsner and L. Wernisch. Markov chains for linear extensions, the two-dimensional case. In *SODA*, pages 239–247, 1997.
- [18] W. Feng, H. Guo, and Y. Yin. Perfect sampling from spatial mixing. *Random Structures & Algorithms*, 61(4):678–709, 2022.
- [19] W. Feng and Y. Yin. On local distributed sampling and counting. In *Proceedings of the 2018 ACM Symposium on Principles of Distributed Computing*, pages 189–198, 2018.
- [20] R. Fernández, A. Procacci, and B. Scoppola. The analyticity region of the hard sphere gas. Improved bounds. *J. Stat. Phys.*, 5:1139–1143, 2007.
- [21] P. A. Ferrari, R. Fernández, and N. L. Garcia. Perfect simulation for interacting point processes, loss networks and Ising models. *Stochastic Processes and their Applications*, 102(1):63–88, 2002.
- [22] D. Gamarnik, D. Katz, and S. Misra. Strong spatial mixing of list coloring of graphs. *Random Structures & Algorithms*, 46(4):599–613, 2015.
- [23] N. L. Garcia. Perfect simulation of spatial processes. *Resenhas do Instituto de Matemática e Estatística da Universidade de São Paulo*, 4(3):283–325, 2000.
- [24] J. Groeneveld. Two theorems on classical many-particle systems. *Phys. Letters*, 3, 1962.
- [25] H. Guo and M. Jerrum. Perfect simulation of the hard disks model by partial rejection sampling. *Annales de l’Institut Henri Poincaré D*, 8(2):159–177, 2021.
- [26] H. Guo, M. Jerrum, and J. Liu. Uniform sampling through the Lovász local lemma. *Journal of the ACM (JACM)*, 66(3):1–31, 2019.
- [27] O. Häggström and K. Nelander. Exact sampling from anti-monotone systems. *Statistica Neerlandica*, 52(3):360–380, 1998.
- [28] O. Häggström, M.-C. N. Van Lieshout, and J. Møller. Characterization results and Markov chain Monte Carlo algorithms including exact simulation for some spatial point processes. *Bernoulli*, 5(4):641–658, 1999.
- [29] T. P. Hayes and C. Moore. Lower bounds on the critical density in the hard disk model via optimized metrics. *arXiv preprint arXiv:1407.1930*, 2014.

- [30] K. He, X. Sun, and K. Wu. Perfect sampling for (atomic) Lovász Local Lemma. *arXiv preprint arXiv:2107.03932*, 2021.
- [31] K. He, C. Wang, and Y. Yin. Sampling Lovasz Local Lemma for general constraint satisfaction solutions in near-linear time, 2022.
- [32] K. He, K. Wu, and K. Yang. Improved bounds for sampling solutions of random CNF formulas, 2022.
- [33] T. Helmuth, W. Perkins, and S. Petti. Correlation decay for hard spheres via Markov chains. *The Annals of Applied Probability*, 32(3):2063–2082, 2022.
- [34] C. Hofer-Temmel and P. Houdebert. Disagreement percolation for Gibbs ball models. *Stochastic Processes and their Applications*, 129(10):3922–3940, 2019.
- [35] R. Holley. Possible rates of convergence in finite range, attractive spin systems. *Part. Syst. Random Media Large Deviat.*, 41:215, 1985.
- [36] M. Huber. Spatial birth–death swap chains. *Bernoulli*, 18(3):1031–1041, 2012.
- [37] M. Huber. Nearly optimal Bernoulli factories for linear functions. *Combin. Probab. Comput.*, 25(4):577–591, 2016.
- [38] M. Huber, E. Vilella, D. Rozenfeld, and J. Xu. Bounds on the artificial phase transition for perfect simulation of hard core Gibbs processes. *Involve, a Journal of Mathematics*, 5(3):247–255, 2013.
- [39] M. Isobe. Hard sphere simulation in statistical physics—methodologies and applications. *Molecular Simulation*, 42(16):1317–1329, 2016.
- [40] V. Jain, A. Sah, and M. Sawhney. Perfectly sampling $k \geq (8/3 + o(1))\Delta$ -colorings in graphs. In *Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing*, pages 1589–1600, 2021.
- [41] M. Jenssen, M. Michelen, and M. Ravichandran. Quasipolynomial-time algorithms for repulsive Gibbs point processes. *arXiv preprint arXiv:2209.10453*, 2022.
- [42] M. Jerrum and A. Sinclair. The Markov chain Monte Carlo method: an approach to approximate counting and integration. *Approximation algorithms for NP-hard problems*, pages 482–520, 1996.
- [43] R. Kannan, M. W. Mahoney, and R. Montenegro. Rapid mixing of several Markov chains for a hard-core model. In *Algorithms and computation*, volume 2906 of *Lecture Notes in Comput. Sci.*, pages 663–675. Springer, Berlin, 2003.
- [44] F. P. Kelly and B. D. Ripley. A note on Strauss’s model for clustering. *Biometrika*, pages 357–360, 1976.

- [45] W. S. Kendall. Perfect simulation for the area-interaction point process. In *Probability towards 2000*, pages 218–234. Springer, 1998.
- [46] W. S. Kendall and J. Møller. Perfect simulation using dominating processes on ordered spaces, with application to locally stable point processes. *Advances in Applied Probability*, pages 844–865, 2000.
- [47] T. Kötzing. Concentration of first hitting times under additive drift. In *Proceedings of the 2014 Annual Conference on Genetic and Evolutionary Computation*, pages 1391–1398, 2014.
- [48] J. Lengler. Drift analysis. In *Theory of Evolutionary Computation*, pages 89–131. Springer, 2020.
- [49] B. Li, Y. Nishikawa, P. Höllmer, L. Carillo, A. Maggs, and W. Krauth. Hard-disk pressure computations—a historic perspective. *The Journal of Chemical Physics*, 157(23):234111, 2022.
- [50] J. Liu, A. Sinclair, and P. Srivastava. Correlation decay and partition function zeros: Algorithms and phase transitions. *SIAM Journal on Computing*, (0):FOCS19–200, 2022.
- [51] L. Lovasz and P. Winkler. Exact mixing in an unknown Markov chain. *The Electronic Journal of Combinatorics*, pages R15–R15, 1995.
- [52] H. Löwen. Fun with hard spheres. In *Statistical physics and spatial statistics*, volume 554, pages 295–331. Springer, 2000.
- [53] P. Lu and Y. Yin. Improved FPTAS for multi-spin systems. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques: 16th International Workshop, APPROX 2013, and 17th International Workshop, RANDOM 2013, Berkeley, CA, USA, August 21-23, 2013. Proceedings*, pages 639–654. Springer, 2013.
- [54] F. Martinelli. Lectures on Glauber dynamics for discrete spin models. *Lectures on probability theory and statistics (Saint-Flour, 1997)*, 1717:93–191, 1999.
- [55] N. Metropolis, A. W. Rosenbluth, M. N. Rosenbluth, A. H. Teller, and E. Teller. Equation of state calculations by fast computing machines. *The Journal of Chemical Physics*, 21(6):1087–1092, 1953.
- [56] M. Michelen and W. Perkins. Potential-weighted connective constants and uniqueness of Gibbs measures. *arXiv preprint arXiv:2109.01094*, 2021.
- [57] M. Michelen and W. Perkins. Analyticity for classical gasses via recursion. *Communications in Mathematical Physics*, pages 1–22, 2022.

- [58] M. Michelen and W. Perkins. Strong spatial mixing for repulsive point processes. *Journal of Statistical Physics*, 189(1):9, 2022.
- [59] M. Mitzenmacher and E. Upfal. *Probability and computing: Randomization and probabilistic techniques in algorithms and data analysis*. Cambridge University Press, 2017.
- [60] S. B. Moka and D. P. Kroese. Perfect sampling for Gibbs point processes using partial rejection sampling. *Bernoulli*, 26(3):2082–2104, 2020.
- [61] J. Møller. A review of perfect simulation in stochastic geometry. *Lecture Notes-Monograph Series*, pages 333–355, 2001.
- [62] J. Møller and R. P. Waagepetersen. *Statistical inference and simulation for spatial point processes*. CRC Press, 2003.
- [63] D. J. Murdoch and P. J. Green. Exact sampling from a continuous state space. *Scandinavian Journal of Statistics*, 25(3):483–502, 1998.
- [64] S. Nacu and Y. Peres. Fast simulation of new coins from old. *The Annals of Applied Probability*, 15(1A):93–115, 2005.
- [65] O. Penrose. Convergence of fugacity expansions for fluids and lattice gases. *Journal of Mathematical Physics*, 4(10):1312–1320, 1963.
- [66] J. G. Propp and D. B. Wilson. Exact sampling with coupled Markov chains and applications to statistical mechanics. *Random Structures & Algorithms*, 9(1-2):223–252, 1996.
- [67] J. G. Propp and D. B. Wilson. How to get a perfectly random sample from a generic Markov chain and generate a random spanning tree of a directed graph. *Journal of Algorithms*, 27(2):170–217, 1998.
- [68] D. Randall. Rapidly mixing Markov chains with applications in computer science and physics. *Computing in Science & Engineering*, 8(2):30–41, 2006.
- [69] G. Regts. Absence of zeros implies strong spatial mixing. *Probability Theory and Related Fields*, pages 1–21, 2023.
- [70] D. Ruelle. Correlation functions of classical gases. *Annals of Physics*, 25:109–120, 1963.
- [71] D. Ruelle. *Statistical mechanics: Rigorous results*. World Scientific, 1999.
- [72] A. Sinclair, P. Srivastava, D. Štefankovič, and Y. Yin. Spatial mixing and the connective constant: Optimal bounds. *Probability Theory and Related Fields*, 168(1-2):153–197, 2017.

- [73] Y. Spinka. Finitary codings for spatial mixing Markov random fields. *Ann. Probab.*, 48(3):1557–1591, 2020.
- [74] D. J. Strauss. A model for clustering. *Biometrika*, 62(2):467–475, 1975.
- [75] D. W. Stroock and B. Zegarlinski. The logarithmic Sobolev inequality for discrete spin systems on a lattice. *Communications in Mathematical Physics*, 149(1):175–193, 1992.
- [76] M. Van Lieshout. *Markov point processes and their applications*. World Scientific, 2000.
- [77] D. Weitz. Counting independent sets up to the tree threshold. In *Proceedings of the Thirty-Eighth Annual ACM Symposium on Theory of Computing*, STOC 2006, pages 140–149. ACM, 2006.
- [78] W. Wood, F. Parker, and J. Jacobson. Recent Monte Carlo calculations of the equation of state of Lenard–Jones and hard sphere molecules. *Il Nuovo Cimento (1955-1965)*, 9:133–143, 1958.

A Measure theory and conditional expectations

A.1 Conditional expectation

We start with a brief recap of the notation used in the appendix.

Let $(\Omega, \mathcal{A}, \mathbb{P})$ be a probability space. For an event $A \in \mathcal{A}$ with $\mathbb{P}[A] > 0$, we write \mathbb{P}_A for the probability measure $\mathbb{P}_A[\cdot] = \mathbb{P}[\cdot \mid A]$ on (Ω, \mathcal{A}) . Note that for all events $A, B \in \mathcal{A}$ with $\mathbb{P}[A \cap B] > 0$ it holds that $\mathbb{P}_{A \cap B} = (\mathbb{P}_A)_B$. Let $f, g : \Omega \rightarrow \mathbb{R}$ be measurable maps. We denote by $\mathbb{E}[f]$ the expectation of f under the measure \mathbb{P} . For a sub- σ -field $\mathcal{F} \subseteq \mathcal{A}$, we write $\mathbb{E}[f \mid \mathcal{F}]$ as a placeholder for any version of a conditional expectation of f given \mathcal{F} under the probability measure \mathbb{P} . Further, we write $\mathbb{E}[\cdot \mid f]$ for conditional expectations given $\sigma(f)$, the σ -field generated by f , and $\mathbb{E}[\cdot \mid f, g]$ for conditional expectations given $\sigma(\sigma(f) \cup \sigma(g))$. Finally, for an event $A \in \mathcal{A}$ with $\mathbb{P}[A] > 0$ and a sub- σ -field $\mathcal{F} \subseteq \mathcal{A}$, we write $\mathbb{E}[f \mid A]$ for the expectation of f under \mathbb{P}_A and $\mathbb{E}[f \mid \mathcal{F}; A]$ for the conditional expectation of f given \mathcal{F} under the measure \mathbb{P}_A .

The following two statements allow us to relate conditional expectations under different probability distributions.

Lemma A.1. *Let $(\Omega, \mathcal{A}, \mathbb{P})$ be a probability space, X be an integrable random variable, let $\mathcal{F} \subseteq \mathcal{A}$ be a sub- σ -field and let $A \in \mathcal{A}$ with $\mathbb{P}[A] > 0$. Then*

$$\mathbb{E}[\mathbb{1}_A X \mid \mathcal{F}] = \mathbb{E}[X \mid \mathcal{F}; A] \mathbb{E}[\mathbb{1}_A \mid \mathcal{F}]$$

\mathbb{P} -almost surely. Moreover, it holds that

$$\mathbb{E}[X \mid \mathcal{F}; A] = \frac{\mathbb{E}[\mathbb{1}_A X \mid \mathcal{F}]}{\mathbb{E}[\mathbb{1}_A \mid \mathcal{F}]}$$

\mathbb{P}_A -almost surely.

Proof. By definition, $\mathbb{E}[X \mid \mathcal{F}; A]\mathbb{E}[\mathbf{1}_A \mid \mathcal{F}]$ is \mathcal{F} -measurable. Moreover, for any $B \in \mathcal{F}$ it holds that

$$\begin{aligned}\mathbb{E}[\mathbf{1}_B \mathbb{E}[X \mid \mathcal{F}; A]\mathbb{E}[\mathbf{1}_A \mid \mathcal{F}]] &= \mathbb{E}[\mathbf{1}_B \mathbf{1}_A \mathbb{E}[X \mid \mathcal{F}; A]] \\ &= \mathbb{E}[\mathbf{1}_A] \mathbb{E}[\mathbf{1}_B \mathbb{E}[X \mid \mathcal{F}; A] \mid A] \\ &= \mathbb{E}[\mathbf{1}_A] \mathbb{E}[\mathbf{1}_B X \mid A] \\ &= \mathbb{E}[\mathbf{1}_B \mathbf{1}_A X] \\ &= \mathbb{E}[\mathbf{1}_B \mathbb{E}[\mathbf{1}_A X \mid \mathcal{F}]].\end{aligned}$$

Thus, it holds that

$$\mathbb{E}[\mathbf{1}_A X \mid \mathcal{F}] = \mathbb{E}[X \mid \mathcal{F}; A]\mathbb{E}[\mathbf{1}_A \mid \mathcal{F}]$$

\mathbb{P} -almost surely. Next, observe that for $A \in \mathcal{A}$ with $\mathbb{P}[A] > 0$ it holds that $\mathbb{E}[A \mid \mathcal{F}] > 0$ \mathbb{P}_A -almost surely. Thus, it follows immediately that

$$\mathbb{E}[X \mid \mathcal{F}; A] = \frac{\mathbb{E}[\mathbf{1}_A X \mid \mathcal{F}]}{\mathbb{E}[\mathbf{1}_A \mid \mathcal{F}]}$$

\mathbb{P}_A -almost surely. □

The following properties can be concluded.

Lemma A.2. *Let $(\Omega, \mathcal{A}, \mathbb{P})$ be a probability space, X be an integrable random variable and $\mathcal{F} \subseteq \mathcal{A}$ be a sub- σ -field. Let $A_1, \dots, A_n \in \mathcal{A}$ be disjoint and such that $\mathbb{P}[A_i] > 0$ for all $1 \leq i \leq n$ and $\mathbb{P}[\bigcup_{i=1}^n A_i] = 1$. If an \mathcal{F} -measurable function $f : \Omega \rightarrow \mathbb{R}$ is a version of $\mathbb{E}[X \mid \mathcal{F}; A_i]$ for all $1 \leq i \leq n$, then f is also a version of $\mathbb{E}[X \mid \mathcal{F}]$.*

Proof. Since the events $A_1, \dots, A_n \in \mathcal{A}$ and satisfy $\mathbb{P}[\bigcup_{i=1}^n A_i] = 1$, we have

$$X = X \mathbf{1}_{\bigcup_{i=1}^n A_i} = \sum_{i=1}^n X \mathbf{1}_{A_i}$$

\mathbb{P} -almost surely. Thus, by linearity of expectation, we have

$$\mathbb{E}[X \mid \mathcal{F}] = \sum_{i=1}^n \mathbb{E}[X \mathbf{1}_{A_i} \mid \mathcal{F}].$$

Furthermore, since $\mathbb{P}[A_i] > 0$ for all $1 \leq i \leq n$ Lemma A.1 gives

$$\sum_{i=1}^n \mathbb{E}[X \mathbf{1}_{A_i} \mid \mathcal{F}] = \sum_{i=1}^n \mathbb{E}[X \mid \mathcal{F}; A_i]\mathbb{E}[\mathbf{1}_{A_i} \mid \mathcal{F}] = f \sum_{i=1}^n \mathbb{E}[\mathbf{1}_{A_i} \mid \mathcal{F}]$$

\mathbb{P} -almost surely. Finally, observing that

$$\sum_{i=1}^n \mathbb{E}[\mathbf{1}_{A_i} \mid \mathcal{F}] = \mathbb{E}[\mathbf{1}_{\bigcup_{i=1}^n A_i} \mid \mathcal{F}] = 1$$

concludes the proof. □

Lemma A.3. *Let $(\Omega, \mathcal{A}, \mathbb{P})$ be a probability space, X be an integrable random variable, let $\mathcal{F} \subseteq \mathcal{A}$ be a sub- σ -field and let $A \in \mathcal{A}$ with $\mathbb{P}[A] > 0$.*

- (1) *If $A \in \mathcal{F}$ then $\mathbb{E}[X \mid \mathcal{F}] = \mathbb{E}[X \mid \mathcal{F}; A]$ \mathbb{P}_A -almost surely.*
- (2) *If $\mathbb{E}[\mathbf{1}_A \mid \mathcal{F}] = \mathbb{E}[\mathbf{1}_A \mid \mathcal{F}, \mathcal{G}]$ for a σ -field $\mathcal{G} \subseteq \mathcal{A}$ with $\sigma(X) \subseteq \mathcal{G}$, then $\mathbb{E}[X \mid \mathcal{F}] = \mathbb{E}[X \mid \mathcal{F}; A]$ \mathbb{P}_A -almost surely.*

Proof. For (1), note that for $A \in \mathcal{F}$ we have $\mathbb{E}[X\mathbf{1}_A \mid \mathcal{F}] = \mathbf{1}_A\mathbb{E}[X \mid \mathcal{F}]$. Moreover, Lemma A.1 yields

$$\mathbb{E}[X\mathbf{1}_A \mid \mathcal{F}] = \mathbb{E}[\mathbf{1}_A \mid \mathcal{F}]\mathbb{E}[X \mid \mathcal{F}; A] = \mathbf{1}_A\mathbb{E}[X \mid \mathcal{F}; A].$$

Thus, we have $\mathbf{1}_A\mathbb{E}[X \mid \mathcal{F}] = \mathbf{1}_A\mathbb{E}[X \mid \mathcal{F}; A]$ and, in particular, $\mathbb{E}[X \mid \mathcal{F}] = \mathbb{E}[X \mid \mathcal{F}; A]$ \mathbb{P}_A -almost surely.

For (2), observe that

$$\begin{aligned} \mathbb{E}[X\mathbf{1}_A \mid \mathcal{F}] &= \mathbb{E}[\mathbb{E}[X\mathbf{1}_A \mid \mathcal{F}, \mathcal{G}] \mid \mathcal{F}] \\ &= \mathbb{E}[X\mathbb{E}[\mathbf{1}_A \mid \mathcal{F}, \mathcal{G}] \mid \mathcal{F}] \\ &= \mathbb{E}[X\mathbb{E}[\mathbf{1}_A \mid \mathcal{F}] \mid \mathcal{F}] \\ &= \mathbb{E}[X \mid \mathcal{F}]\mathbb{E}[\mathbf{1}_A \mid \mathcal{F}] \end{aligned}$$

\mathbb{P} -almost surely, where the second equality follows from $\sigma(X) \subseteq \mathcal{G}$ and the third follows from $\mathbb{E}[\mathbf{1}_A \mid \mathcal{F}] = \mathbb{E}[\mathbf{1}_A \mid \mathcal{F}, \mathcal{G}]$. The claim now follows from Lemma A.1. □

A.2 Regular conditional distributions

Consider a probability space $(\Omega, \mathcal{A}, \mathbb{P})$ with a sub- σ -field $\mathcal{F} \subseteq \mathcal{A}$, a measure space (D, \mathcal{D}) and (D, \mathcal{D}) -valued random variable X . A map $Q : \Omega \times \mathcal{D} \rightarrow [0, 1]$ is called a regular conditional distribution of X given \mathcal{F} if

- (1) $Q(\omega, \cdot)$ is a probability measure on (D, \mathcal{D}) for all $\omega \in \Omega$ and
- (2) $Q(\cdot, A)$ is a version of $\mathbb{E}[\mathbf{1}_{X \in A} \mid \mathcal{F}]$ for all $A \in \mathcal{D}$.

The following statements makes regular conditional distributions particularly useful.

Theorem A.4 ([11][Theorem 2.19]). *Let $(\Omega, \mathcal{A}, \mathbb{P})$ be a probability space, let X be a (D, \mathcal{D}) -valued random variable and let Y be a (E, \mathcal{E}) -valued random variable. If $Q : \Omega \times \mathcal{D}$ is a regular conditional distribution of X given $\sigma(Y)$, then, for all $\mathcal{D} \otimes \mathcal{E}$ -measurable $f : D \times E \rightarrow \mathbb{R}_{\geq 0}$ it holds that*

$$\mathbb{E}[f(X, Y) \mid Y](\omega) = \int_D f(x, Y(\omega))Q(\omega, dx)$$

for \mathbb{P} -almost all $\omega \in \Omega$.

Moreover, the following lemma helps to identify regular conditional distributions based on a π -system.

Lemma A.5. *Let $(\Omega, \mathcal{A}, \mathbb{P})$ be a probability space and let $\mathcal{F} \subseteq \mathcal{A}$ be a sub- σ -field. Let X be a (D, \mathcal{D}) -valued random variable on $(\Omega, \mathcal{A}, \mathbb{P})$ and let $\mathcal{G} \subseteq \mathcal{D}$ be a π -system that generates \mathcal{D} . Assume there is a function $Q : \Omega \times \mathcal{D}$ such that*

- (1) $Q(\omega, \cdot)$ is a probability distribution on (D, \mathcal{D}) for all $\omega \in \Omega$ and
- (2) $Q(\cdot, G)$ is a version of $\mathbb{E}[\mathbb{1}_{X \in G} \mid \mathcal{F}]$ for all $G \in \mathcal{G}$.

Then $Q(\cdot, A)$ is a version of $\mathbb{E}[\mathbb{1}_{X \in A} \mid \mathcal{F}]$ for all $A \in \mathcal{D}$ and, in particular, Q is a regular conditional distribution for X given \mathcal{F} .

Proof. Consider the set of events

$$\mathcal{H} = \{A \in \mathcal{D} \mid Q(\cdot, A) \text{ is a version of } \mathbb{E}[\mathbb{1}_{X \in A} \mid \mathcal{F}]\}.$$

Our goal is to prove $\mathcal{H} = \mathcal{D}$. To this end, note that $\mathcal{G} \subseteq \mathcal{H} \subseteq \mathcal{D}$. Thus, if we prove that \mathcal{H} is a Dynkin system, then the π - λ Theorem implies that $\mathcal{D} = \sigma(\mathcal{G}) \subseteq \mathcal{H}$, which proves our claim. To show that \mathcal{H} is a Dynkin system, we need to argue that $D \in \mathcal{H}$ and that \mathcal{H} is closed under complements and countable disjoint unions.

To see that $D \in \mathcal{H}$, note that $Q(\omega, D) = 1$ for all $\omega \in \Omega$. Thus, $Q(\cdot, D)$ is trivially \mathcal{F} -measurable. Moreover, for any $B \in \mathcal{F}$, it holds that

$$\mathbb{E}[\mathbb{1}_B Q(\cdot, D) \mid \mathcal{F}] = \mathbb{E}[\mathbb{1}_B \mid \mathcal{F}] = \mathbb{E}[\mathbb{1}_B \mathbb{1}_{X \in D} \mid \mathcal{F}],$$

which shows that $Q(\cdot, D)$ is indeed a version of $\mathbb{E}[\mathbb{1}_{X \in D} \mid \mathcal{F}]$.

Next, fix some $G \in \mathcal{H}$ and observe that $Q(\omega, (G)^c) = 1 - Q(\omega, G)$ for all $\omega \in \Omega$. Since $Q(\cdot, G)$ is by assumption a version of $\mathbb{E}[\mathbb{1}_{X \in G} \mid \mathcal{F}]$ (therefore \mathcal{F} -measurable), this shows that $Q(\cdot, (G)^c)$ is \mathcal{F} -measurable. Moreover, for all $B \in \mathcal{F}$, we have

$$\begin{aligned} \mathbb{E}[\mathbb{1}_B Q(\cdot, (G)^c) \mid \mathcal{F}] &= \mathbb{E}[\mathbb{1}_B \mid \mathcal{F}] - \mathbb{E}[\mathbb{1}_B Q(\cdot, G) \mid \mathcal{F}] \\ &= \mathbb{E}[\mathbb{1}_B \mathbb{1}_{X \in D} \mid \mathcal{F}] - \mathbb{E}[\mathbb{1}_B \mathbb{1}_{X \in G} \mid \mathcal{F}] \\ &= \mathbb{E}[\mathbb{1}_B \mathbb{1}_{X \in (G)^c} \mid \mathcal{F}], \end{aligned}$$

which proves that $(G)^c \in \mathcal{H}$.

Finally, consider some sequence of disjoint events $(G_n)_{n \in \mathbb{N}} \in \mathcal{H}^{\mathbb{N}}$ and set $G = \bigcup_{n \in \mathbb{N}} G_n$. Note that $Q(\omega, G) = \sum_{n \in \mathbb{N}} Q(\omega, G_n)$ for all $\omega \in \Omega$. Since each function $Q(\cdot, G_n)$ is a version of $\mathbb{E}[\mathbb{1}_{X \in G_n} \mid \mathcal{F}]$, this implies that $Q(\cdot, G)$ is \mathcal{F} -measurable. Moreover, for all $B \in \mathcal{F}$ it holds that

$$\mathbb{E}[\mathbb{1}_B Q(\cdot, G) \mid \mathcal{F}] = \sum_{n \in \mathbb{N}} \mathbb{E}[\mathbb{1}_B Q(\cdot, G_n) \mid \mathcal{F}] = \sum_{n \in \mathbb{N}} \mathbb{E}[\mathbb{1}_B \mathbb{1}_{X \in G_n} \mid \mathcal{F}] = \mathbb{E}[\mathbb{1}_B \mathbb{1}_{X \in G} \mid \mathcal{F}],$$

showing that $G \in \mathcal{H}$. Thus, \mathcal{H} is a Dynkin system, which concludes the proof. \square

B Hitting times and tail bounds

We frequently make use of the following version of Wald's identity.

Lemma B.1. *Let $(\Omega, \mathcal{A}, \mathbb{P})$ be a probability space, let $(X_n)_{n \in \mathbb{N}_{\geq 1}}$ be a sequence of random variables on $(\Omega, \mathcal{A}, \mathbb{P})$ with values in $\mathbb{R}_{\geq 0}$ and let $\mathcal{F} \subseteq \mathcal{A}$ be a sub- σ -field. Suppose there is a \mathcal{F} -measurable random variable M such that for all $n \in \mathbb{N}$ it holds that $\mathbb{E}[X_n \mid \mathcal{F}] \leq M$ almost surely. Let N be a random variable in \mathbb{N} such that for all $n \in \mathbb{N}$ it holds that $\mathbb{E}[X_n \mathbf{1}_{N \geq n} \mid \mathcal{F}] = \mathbb{E}[X_n \mid \mathcal{F}] \mathbb{E}[\mathbf{1}_{N \geq n} \mid \mathcal{F}]$ almost surely. Then $\mathbb{E}\left[\sum_{n=1}^N X_n \mid \mathcal{F}\right] \leq M \mathbb{E}[N \mid \mathcal{F}]$ almost surely.*

Proof. Using monotone convergence we have

$$\begin{aligned} \mathbb{E}\left[\sum_{n=1}^N X_n \mid \mathcal{F}\right] &= \sum_{n=1}^{\infty} \mathbb{E}[\mathbf{1}_{N \geq n} X_n \mid \mathcal{F}] = \sum_{n=1}^{\infty} \mathbb{E}[X_n \mid \mathcal{F}] \mathbb{E}[\mathbf{1}_{N \geq n} \mid \mathcal{F}] \\ &\leq M \mathbb{E}\left[\sum_{n=1}^{\infty} \mathbf{1}_{N \geq n} \mid \mathcal{F}\right] = M \mathbb{E}[N \mid \mathcal{F}] \quad \square \end{aligned}$$

Moreover, we use the following drift theorem to bound the expected number of iterations of our sampling algorithm.

Theorem B.2 ([48, Theorem 1], [47, Theorem 2]). *Let $(X_t)_{t \in \mathbb{N}_0}$ be an integrable random process over \mathbb{R} that is adapted to a filtration $(\mathcal{F}_t)_{t \in \mathbb{N}_0}$ and let $T = \inf\{t \in \mathbb{N}_0 \mid X_t \leq 0\}$. Assume*

- a) $X_t \mathbf{1}_{T \geq t} \geq 0$ for all $t \in \mathbb{N}$ and
- b) there is some $\varepsilon \in \mathbb{R}_{>0}$ such that $\mathbb{E}[(X_t - X_{t+1}) \mathbf{1}_{T > t} \mid \mathcal{F}_t] \geq \varepsilon \mathbf{1}_{T > t}$ for all $t \in \mathbb{N}$.

Then $\mathbb{E}[T] \leq \frac{\mathbb{E}[X_0]}{\varepsilon}$. Further, suppose that

- c) $X_0 \leq x$ for some $x \in \mathbb{R}_{>0}$ and
- d) there is some $c \in \mathbb{R}_{>0}$ such that $|X_t - X_{t+1}| < c$ for all $t \in \mathbb{N}_0$.

Then, for all $s \geq \frac{2x}{\varepsilon}$, $\mathbb{P}[T \geq s] \leq \exp\left(-\frac{s\varepsilon^2}{16c^2}\right)$.

Finally, we make use of the following tail bound for Poisson random variables.

Theorem B.3 ([59, Theorem 5.4]). *Let $Y \sim \text{Pois}(\rho)$ for some $\rho \in \mathbb{R}_{>0}$. For all $y > \rho$ it holds that $\mathbb{P}[Y \geq y] \leq e^{-\rho} \left(\frac{e\rho}{y}\right)^y$.*

In particular, we use the following corollary of the above bound.

Corollary B.4. *Let $Y \sim \text{Pois}(\rho)$ for some $\rho \in \mathbb{R}_{>0}$. For all $\gamma > 1$ all $y \geq e^\gamma \rho$ it holds that $\mathbb{P}[Y \geq y] \leq e^{-(\gamma-1)y}$.*

Proof. Since $y \geq e^\gamma \rho > \rho$, Theorem B.3 implies that

$$\mathbb{P}[Y \geq y] \leq e^{-\rho} \left(\frac{e\rho}{y}\right)^y \leq \left(\frac{e\rho}{e^\gamma \rho}\right)^y = e^{-(\gamma-1)y}. \quad \square$$

C Gibbs point processes

Here we collect some useful lemmas about Gibbs point processes.

The following technical lemma will come in handy.

Lemma C.1. *Let $\Lambda'_1 \subset \Lambda' \subset \mathbb{R}^d$ be bounded and measurable. Denote by $\pi_{\Lambda', \Lambda'_1} : \mathcal{N}_{\Lambda'} \rightarrow \mathcal{N}_{\Lambda'_1}$ the projection $\eta \mapsto \eta \cap \Lambda'_1$. Then $\pi_{\Lambda', \Lambda'_1}$ is $\mathfrak{R}_{\Lambda'}$ -measurable. Moreover, for $\Lambda'_2 = \Lambda' \setminus \Lambda'_1$ it holds that*

$$\mathcal{G} = \left\{ \pi_{\Lambda', \Lambda'_1}^{-1}(A_1) \cap \pi_{\Lambda', \Lambda'_2}^{-1}(A_2) \mid A_1 \in \mathfrak{R}_{\Lambda'_1}, A_2 \in \mathfrak{R}_{\Lambda'_2} \right\}$$

is a π -system that generates $\mathfrak{R}_{\Lambda'}$.

Proof. We start by showing that $\pi_{\Lambda', \Lambda'_1}$ is $\mathfrak{R}_{\Lambda'}$ -measurable. To this end, set $E_{B,k}^{(1)} = \{\eta \in \mathcal{N}_{\Lambda'_1} \mid |\eta \cap B| = k\} \in \mathfrak{R}_{\Lambda'_1}$ for every Borel set $B \subseteq \Lambda'_1$ and $k \in \mathbb{N}_0$, and observe that $\mathcal{E}^{(1)} = \{E_{B,k} \mid k \in \mathbb{N}_0, B \subseteq \Lambda'_1 \text{ Borel}\}$ generates $\mathfrak{R}_{\Lambda'_1}$. Thus, it suffices to show for every $E_{B,k}^{(1)} \in \mathcal{E}^{(1)}$ that $\pi_{\Lambda', \Lambda'_1}^{-1}(E_{B,k}^{(1)}) \in \mathfrak{R}_{\Lambda'}$. By definition, we have

$$\pi_{\Lambda', \Lambda'_1}^{-1}(E_{B,k}^{(1)}) = \{\eta \in \mathcal{N}_{\Lambda'} \mid |(\eta \cap \Lambda'_1) \cap B| = k\} = \{\eta \in \mathcal{N}_{\Lambda'} \mid |\eta \cap B| = k\} \in \mathfrak{R}_{\Lambda'},$$

which proves the first part of the claim.

For the second part, note that \mathcal{G} is non-empty and, since $\mathfrak{R}_{\Lambda'_1}$ and $\mathfrak{R}_{\Lambda'_2}$ are closed under countable intersections, the same holds for \mathcal{G} . In particular, \mathcal{G} satisfies the definition of a π -system. It remains to show that \mathcal{G} generates $\mathfrak{R}_{\Lambda'}$. Analogously to the first part of the proof, set $E_{B,k}^{(2)} = \{\eta \in \mathcal{N}_{\Lambda'_2} \mid |\eta \cap B| = k\} \in \mathfrak{R}_{\Lambda'_2}$ for every Borel set $B \subseteq \Lambda'_2$ and $k \in \mathbb{N}_0$, set $E_{B,k} = \{\eta \in \mathcal{N}_{\Lambda'} \mid |\eta \cap B| = k\} \in \mathfrak{R}_{\Lambda'}$ for every Borel set $B \subseteq \Lambda'$ and $k \in \mathbb{N}_0$, and define $\mathcal{E} = \{E_{B,k} \mid k \in \mathbb{N}_0, B \subseteq \Lambda' \text{ Borel}\}$. Since \mathcal{E} generates $\mathfrak{R}_{\Lambda'}$, it suffices to show that $\mathcal{E} \subseteq \sigma(\mathcal{G})$. To this end, fix some $E_{B,k} \in \mathcal{E}$ and observe that

$$\begin{aligned} E_{B,k} &= \bigcup_{k_1=0}^k (\{\eta \in \mathcal{N}_{\Lambda'} \mid |\eta \cap (B \cap \Lambda'_1)| = k_1\} \cap \{\eta \in \mathcal{N}_{\Lambda'} \mid |\eta \cap (B \cap \Lambda'_2)| = k - k_1\}) \\ &= \bigcup_{k_1=0}^k \left(\pi_{\Lambda', \Lambda'_1}^{-1}(E_{B \cap \Lambda'_1, k_1}^{(1)}) \cap \pi_{\Lambda', \Lambda'_2}^{-1}(E_{B \cap \Lambda'_2, k-k_1}^{(2)}) \right). \end{aligned}$$

Thus, every $E_{B,k} \in \mathcal{E}$ can be written as a finite union of sets in \mathcal{G} and $\mathcal{E} \subseteq \sigma(\mathcal{G})$, which proves the claim. \square

In this bounded-range setting, the following lemma might be seen as a version of the spatial Markov property for partition functions.

Lemma C.2. *Let ϕ be a repulsive potential with bounded range $r \in \mathbb{R}_{\geq 0}$. Moreover, let $\Lambda' \subseteq \Lambda$ and let be any activity function $\lambda : \mathbb{R}^d \rightarrow \mathbb{R}_{\geq 0}$. For every two point configurations $\eta_1, \eta_2 \in \mathcal{N}_{\Lambda}$ with $\text{dist}(\Lambda', \eta_1 \ominus \eta_2) \geq r$, where \ominus denotes the symmetric difference, it holds that $Z_{\Lambda'}(\lambda_{\eta_1}) = Z_{\Lambda'}(\lambda_{\eta_2})$.*

Proof. Since the range of ϕ is bounded by r , it holds that $\lambda_{\eta_1}(x) = \lambda_{\eta_2}(x)$ for all $x \in \Lambda$. Therefore, we have

$$Z_{\Lambda'}(\lambda_{\eta_1}) = \sum_{k \geq 0} \frac{1}{k!} \int_{\Lambda'^k} \lambda_{\eta_1}^{\mathbf{x}} e^{-H(\mathbf{x})} d\mathbf{x} = \sum_{k \geq 0} \frac{1}{k!} \int_{\Lambda'^k} \lambda_{\eta_2}^{\mathbf{x}} e^{-H(\mathbf{x})} d\mathbf{x} = Z_{\Lambda'}(\lambda_{\eta_2}),$$

which proves the claim. □