

REJECTION- AND IMPORTANCE-SAMPLING-BASED PERFECT SIMULATION FOR GIBBS HARD-SPHERE MODELS

SARAT MOKA,* *University of Queensland*

SANDEEP JUNEJA,** *Tata Institute of Fundamental Research*

MICHEL MANDJES,*** *University of Amsterdam*

Abstract

Coupling-from-the-past (CFTP) methods have been used to generate perfect samples from finite Gibbs hard-sphere models, an important class of spatial point processes consisting of a set of spheres with the centers on a bounded region that are distributed as a homogeneous Poisson point process (PPP) conditioned so that spheres do not overlap with each other. We propose an alternative importance-sampling-based rejection methodology for the perfect sampling of these models. We analyze the asymptotic expected running time complexity of the proposed method when the intensity of the reference PPP increases to infinity while the (expected) sphere radius decreases to zero at varying rates. We further compare the performance of the proposed method analytically and numerically with that of a naive rejection algorithm and of popular dominated CFTP algorithms. Our analysis relies upon identifying large deviations decay rates of the non-overlapping probability of spheres whose centers are distributed as a homogeneous PPP.

Keywords: Exact simulation; dominated coupling from the past; large deviations; non-overlapping probability

2010 Mathematics Subject Classification: Primary 65C05

Secondary 60G55; 60F10; 68W40

1. Introduction

Perfect sampling, that is, the generation of unbiased samples from a target distribution (also referred to as perfect simulation or exact sampling), is an important and exciting area of research in stochastic simulation. In this paper, we introduce and investigate a novel methodology for generating perfect samples of finite Gibbs hard-sphere models, which are an important family of Gibbs point processes. Roughly, a Gibbs hard-sphere model can be described as a set of spheres such that their centers constitute a Poisson point process (PPP) on a bounded Euclidean space, conditioned that no two spheres overlap with each other. The proposed methodology combines *importance sampling* (IS) and *acceptance–rejection* (AR) techniques to achieve substantial performance improvement in certain important regimes of interest. In

Received 11 March 2018; revision received 14 December 2020.

* Postal address: School of Mathematics and Physics, Coopers Road, St Lucia, QLD, Australia 4072. Email: s.babumoka@uq.edu.au

** Postal address: Homi Bhabha Road, Colaba, Mumbai, India 400005. Email: juneja@tifr.res.in

*** Postal address: Korteweg-de Vries Institute for Mathematics, University of Amsterdam, Science Park 904, 1098 XH Amsterdam, The Netherlands. Email: m.r.h.mandjes@uva.nl

© The Author(s), 2021. Published by Cambridge University Press on behalf of Applied Probability Trust.

statistical physics, there is a large body of work related to Gibbs hard-sphere models; see, e.g., [1, 2, 6, 24, 28, 30, 35, 37]. These models are important also in modeling adsorption of latexes or proteins on solid surfaces (see [38, 40] and references therein). For the analysis of wireless communication networks, it is common to use Gibbs hard-sphere models to model base-stations in a cellular network, because normally no two base-stations are to be placed closer than a certain distance from each other [18, 39]. Our results can be used to assess the stationary behavior of code-division multiple access wireless networks.

Literature review: The existing literature offers several perfect sampling methods for Gibbs hard-sphere models. Among these, the dominated coupling-from-the-past (CFTP) methods are most prominent; these are based on the seminal paper by Propp and Wilson [36] (see [22, 25–27]). Another well-known perfect sampling method for the Gibbs hard-sphere models is the backward–forward algorithm by Ferrari *et al.* [12]; see also [16, 23]. For some of the applications of perfect sampling for these models, refer to [7, 8, 34]. For other related literature on perfect sampling for spatial point processes, refer to [19, 32].

As mentioned in [16], all the existing methods are, in some sense, complementary to each other. They take advantage of the important property that the distribution of a Gibbs hard-sphere model can be realized as an invariant measure of a spatial birth-and-death process, called the *interaction process*. For example, the main ingredient of the dominated CFTP method is to construct a birth-and-death process backward in time, starting from its steady state at time zero, such that it dominates the interaction process, and then use *thinning* on the dominating process to construct coupled upper and lower bound processes forward in time such that the coalescence of these two bounding processes assures a perfect sample from the target measure, which is the invariant measure of the interaction process. The backward–forward algorithm is based on the construction of the *clan of ancestors*, which uses thinning of a dominating process and extends the applicability to infinite-volume measures. A crucial drawback of the naive AR and the dominated CFTP methods is that they are guaranteed to be efficient only if the intensity of the Gibbs hard-sphere model is close to the intensity of the reference PPP; see [23] for details. In addition, most of the dominated CFTP methods suffer from the so-called *impatient-user bias* (a bias that is induced when a user aborts long runs of the algorithm); see [13, 14, 41].

Our contributions: AR methods are free of the impatient-user bias and involve neither thinning nor coupling (which are crucial for the other methods). Despite being an obvious alternative to the existing methods, to the best of our knowledge, the use of AR methods is still largely unexplored in the context of Gibbs point processes (except for brief discussions, e.g. in [15, 23]). AR methods for Gibbs hard-sphere models are amenable to further algorithmic enhancements that may substantially decrease the expected running time of the algorithm. The methodology proposed here provides one such enhancement. To highlight the significance of the proposed methodology, we compare its running time complexity with that of both the naive AR and the dominated CFTP methods. This effectiveness analysis is based on our large deviations analysis of the *non-overlapping probability*. A brief summary of our results is as follows:

- Our first key contribution is that we conduct a large deviations analysis of the probability of spheres not overlapping with each other when their centers constitute a homogeneous PPP. More specifically, we consider a homogeneous *marked* PPP on $[0, 1]^d$ with intensity λ , where the points are the centers of spheres with independently and identically distributed (i.i.d.) radii as marks, which are independent of the centers and identical in distribution to R/λ^η for a positive bounded random variable R and a constant $\eta > 0$. We

establish large deviations of the probability that spheres do not overlap with each other, as $\lambda \nearrow \infty$. This analysis is useful in the study of the asymptotic behavior of the expected running time complexities of the proposed and the existing perfect sampling methods for the Gibbs hard-sphere models. This analysis may also be of independent interest.

- Our second key contribution is that we propose a novel IS-based AR algorithm for generating perfect samples of the Gibbs hard-sphere model obtained by considering the homogeneous marked PPP conditioned on no overlap of the spheres. This is achieved by partitioning the underlying configuration space and arriving at an appropriate change of measure on each partition. The applicability of the proposed algorithm is illustrated in two scenarios. In the first scenario, all the spheres are assumed to be of a fixed size (i.e., R is a fixed positive constant). We develop a grid-based IS technique under which spheres are generated sequentially so that the chance of spheres overlapping is small, and the corresponding likelihood ratio has a better deterministic upper bound that improves the acceptance probability in each iteration of the algorithm. In the second scenario, we consider the general case where spheres have i.i.d. radii. In this scenario, we divide the underlying configuration space into two sets. On one set, the sum of the volumes of spheres is bounded from below, and on the other set, the volume sum takes small values so that the set consists of only rare configurations. For the first set, we develop a grid-based IS method that is similar to the one described above, and for the second set, we use an *exponential twisting* on the sphere volume distribution. In both the scenarios, the new method provably substantially improves the performance of the algorithm compared to the naive AR method.
- We analytically and numerically compare the performance of the proposed IS-based AR method with that of some of the dominated CFTP methods. The numerical results support our analytical conclusions that the proposed method is substantially more efficient than the existing methods over the *high-density regime*, where $\eta d \leq 1$ and λ is large.

Organization: Section 2 provides a definition of the hard-sphere model. Section 3 presents the large deviations of the non-overlapping probability. In Section 4, we first review a naive AR method and analyze its expected running time complexity, then propose and analyze the IS-based AR method. Section 5 gives a review of the well-known dominated CFTP methods for the hard-sphere models. Section 6 illustrates the efficiency of the proposed methodology using numerical experiments. Section 7 provides a brief conclusion. All proofs are presented in Appendix A.

2. Preliminaries

First we introduce some notation. $X \sim F$ denotes that the distribution of a random object X is F . $\text{Poi}(\lambda)$ and $\text{Bern}(p)$ denote, respectively, Poisson distribution with mean $\lambda > 0$ and Bernoulli distribution with success probability p . The uniform distribution on $[0, 1]$ is denoted by $\text{Unif}(0, 1)$. For an event A , the indicator function $I(A)$ is equal to 1 if A occurs, and to 0 otherwise. A measure μ_1 is absolutely continuous with respect to a measure μ_2 on a measurable set A if $\mu_1(B \cap A) = 0$ for any measurable B such that $\mu_2(B \cap A) = 0$. For any probability measure μ , $\mathbb{P}_\mu(A)$ denotes the probability of an event A under μ , and $\mathbb{E}_\mu[\cdot]$ denotes the associated expectation. We drop the subscript μ when it is not relevant. For any non-negative real-valued functions f and g , we write $f(x) = O(g(x))$ if $\limsup_{x \rightarrow \infty} f(x)/g(x) \leq c$ for some constant $c > 0$, $f(x) = \Omega(g(x))$ if $g(x) = O(f(x))$, and $f(x) = o(g(x))$ if $\limsup_{x \rightarrow \infty} f(x)/g(x) = 0$. We write

$f(x) = \Theta(g(x))$ if both $f(x) = O(g(x))$ and $f(x) = \Omega(g(x))$. For any real number x , the largest integer n such that $n \leq x$ is denoted by $\lfloor x \rfloor$, and the smallest integer n such that $n \geq x$ is denoted by $\lceil x \rceil$. The set of all the non-negative integers is denoted by \mathbb{N}_0 .

A random finite subset $\mathbf{X} = \{X_1, \dots, X_N\}$ of an observation window $W \subset \mathbb{R}^d$ is called a *Poisson point process* (PPP) with a finite intensity measure ν on W if $N \sim \text{Poi}(\nu(W))$ and for every $n \in \mathbb{N}_0$, conditioned on $N = n$, the points X_1, \dots, X_n are i.i.d. with distribution $\nu(dx)/\nu(W)$. A PPP on $[0, 1]^d$ is called a λ -homogeneous PPP with intensity $\lambda > 0$ if the intensity measure $\nu(dx) = \lambda dx$, where dx is Lebesgue measure on W . To each point X_i of the λ -homogeneous PPP on $[0, 1]^d$ we associate a mark, which is a non-negative number interpreted as the radius of a sphere centered at X_i . In particular, a λ -homogeneous marked PPP on $[0, 1]^d$ is a PPP on $W = [0, 1]^d \times [0, \infty)$ with the intensity measure $\nu(dx \times dr) = \lambda dx \times F(dr)$, where F is the distribution of each radius. That is, the centers constitute a λ -homogeneous PPP on $[0, 1]^d$ which is independent of the radii, and the radii are i.i.d. with distribution F . A realization of the marked PPP with n points is denoted by $\mathbf{x} = \{(y_1, r_1), \dots, (y_n, r_n)\}$, where $r_i \geq 0$ is the radius of the sphere centered at $y_i \in [0, 1]^d$. Define $\mathcal{G} = \cup_{n \in \mathbb{N}_0} \mathcal{G}_n$, where

$$\mathcal{G}_n = \{ \mathbf{x} = \{(y_1, r_1), \dots, (y_n, r_n)\} : (y_i, r_i) \in [0, 1]^d \times [0, \infty), \text{ for } i = 1, \dots, n \}.$$

Now we define a *Gibbs hard-sphere model*. Suppose that μ^0 is the distribution of a λ -homogeneous marked PPP as defined above, with F being the distribution of R/λ^η for a constant $\eta > 0$ and a non-negative random variable R . Let $\mathcal{A} \subset \mathcal{G}$ be the set of all configurations with no two spheres overlapping with each other. Then the distribution μ of the Gibbs hard-sphere model is absolutely continuous with respect to μ^0 , with the Radon–Nikodym derivative given by

$$\frac{d\mu}{d\mu^0}(\mathbf{x}) = \frac{I(\mathbf{x} \in \mathcal{A})}{\mathcal{P}(\lambda)}, \quad \mathbf{x} \in \mathcal{G}, \tag{2.1}$$

where the normalizing constant $\mathcal{P}(\lambda)$ is the non-overlapping probability given by

$$\mathcal{P}(\lambda) = \mathbb{P}_{\mu^0}(\mathbf{X} \in \mathcal{A}). \tag{2.2}$$

We refer to the Gibbs hard-sphere model as a *torus-hard-sphere model* if the boundary of the underlying space $[0, 1]^d$ is periodic; that is, a sphere $S(x, a)$ centered at $x \in [0, 1]^d$ with radius r is defined by

$$S(x, r) = \left\{ (y_1 \bmod 1, \dots, y_d \bmod 1) : y = (y_1, \dots, y_d) \in \mathbb{R}^d, \ \|x - y\| < r \right\},$$

where $\|\cdot\|$ is the d -dimensional Euclidean norm and **mod** denotes the modulo operation [9]. If the boundary is not periodic, we refer to the model as a *Euclidean-hard-sphere model*.

From now on, the phrase ‘hard-sphere model’ refers to either of these two models, and we assume that R is bounded from above by a constant $\bar{r} > 0$. In particular, if R is a constant, we take $\bar{r} = R$. Furthermore, we assume that $2\bar{r}/\lambda^\eta < 1$ to avoid certain trivial difficulties such as the possibility of a sphere on the torus overlapping with itself.

3. Large deviations results

In this section, we obtain large deviations results for the non-overlapping probability $\mathcal{P}(\lambda)$. We use these results in analyzing the running time complexity of both the naive and the IS-based AR methods. Hereafter, $\gamma = \pi^{d/2}/\Gamma(d/2 + 1)$, where $\Gamma(\cdot)$ is the gamma function. Note

that the volume of a sphere with radius r is given by γr^d . Define $m_1 := \mathbb{E}[(R + \widehat{R})^d]$, where \widehat{R} is independent and identical in distribution to R , and let

$$\gamma' = \begin{cases} \gamma & \text{if } [0, 1]^d \text{ is treated as the torus,} \\ \gamma/2^d & \text{otherwise.} \end{cases} \tag{3.1}$$

Theorem 3.1. *The non-overlapping probability $\mathcal{P}(\lambda)$ satisfies*

$$\begin{aligned} \lim_{\lambda \rightarrow \infty} \mathcal{P}(\lambda) &= \begin{cases} 1 & \text{if } \eta d > 2, \\ \exp\left(-\frac{\gamma m_1}{2}\right) & \text{if } \eta d = 2, \end{cases} \\ \lim_{\lambda \rightarrow \infty} \left[\frac{1}{\lambda^{2-\eta d}} \log \mathcal{P}(\lambda) \right] &= -\frac{\gamma m_1}{2} \text{ if } 1 < \eta d < 2, \\ \lim_{\lambda \rightarrow \infty} \left[\frac{1}{\lambda} \log \mathcal{P}(\lambda) \right] &= -1 \text{ if } 0 < \eta d < 1. \end{aligned}$$

When $\eta d = 1$, the limit

$$\delta := \lim_{\lambda \rightarrow \infty} \left[\frac{1}{\lambda} \log \mathcal{P}(\lambda) \right]$$

exists and $-1 \leq \delta < 0$. Furthermore, $\delta \nearrow 0$ if $\gamma m_1 \searrow 0$, and

$$\delta \leq -\frac{1}{2} \left(1 - \frac{1}{\gamma' \bar{r}^d} \right)^2$$

if $R \equiv \bar{r}$ and $\gamma' \bar{r}^d > 1$. In addition, for the torus-hard-sphere model,

$$\lim_{\lambda \rightarrow \infty} \left[\mathcal{P}(\lambda) \exp\left(\frac{\gamma m_1}{2} \lambda^{2-\eta d}\right) \right] = 1 \text{ if } 5/3 < \eta d < 2.$$

An important and fundamental characteristic of a Gibbs point process is its *intensity*; see, for example, [29] and references therein, as well as [5]. Roughly speaking, the intensity of a Gibbs point process is the expected number of points of the process per unit volume. There is an interesting connection between the regimes considered in Theorem 3 and the asymptotic intensity of the torus-hard-sphere model. To see this, assume that each sphere has a fixed radius \bar{r}/λ^η . Since the underlying space is $[0, 1]^d$, the intensity $\rho(\lambda)$ of the model is exactly equal to the expected total number of points in a realization of the model. Equivalently, we may consider the fraction of the volume $\mathbf{VF}(\lambda)$ occupied by the spheres, given by $\mathbf{VF}(\lambda) = \rho(\lambda) \gamma \bar{r}^d \lambda^{-\eta d}$. For the torus-hard-sphere model, the volume fraction $\mathbf{VF}(\lambda)$ is bounded from above by $\rho^{\max} \gamma$, where ρ^{\max} is the closest packing density, defined by $\rho^{\max} = \lim_{n \rightarrow \infty} N_n / (n + 1)^d$, with N_n being the maximal number of mutually disjoint unit-radius spheres which are included in the hypercube $[-(n + 1/2), (n + 1/2)]^d$; see [29]. Proposition 3 describes the asymptotic behavior of $\mathbf{VF}(\lambda)$ as $\lambda \rightarrow \infty$ for different values of ηd . In particular, the regime with $\eta d > 1$ is a *low-density regime*, while the regime with $\eta d < 1$ is a *high-density regime*. In the high-density

regime, the intensity of the hard-sphere model is much smaller than the intensity λ of the reference PPP.

Proposition 3.1. *For the torus-hard-sphere model with a fixed radius $R = \bar{r}$,*

$$\begin{aligned} \lim_{\lambda \nearrow \infty} \frac{\text{VF}(\lambda)}{\gamma \bar{r}^d \lambda^{1-\eta d}} &= 1 \quad \text{if } \eta d > 1, \\ \lim_{\lambda \nearrow \infty} \text{VF}(\lambda) &= \rho^{\max} \gamma \quad \text{if } \eta d < 1, \\ \lim_{\lambda \nearrow \infty} \text{VF}(\lambda) &< \rho^{\max} \gamma \quad \text{if } \eta d = 1. \end{aligned}$$

4. Acceptance–rejection-based algorithms

In Section 4.1, we present a naive AR algorithm for generating perfect samples of the hard-sphere model and analyze its expected running time complexity. We then proceed to present and analyze our IS-based AR algorithm, where the key idea is to partition the configuration space \mathcal{G} so that a well chosen IS technique can be implemented on each partition. One such IS for the hard-sphere model is the reference IS presented in Section 4.3, where spheres are generated sequentially so that, whenever possible, the center of each sphere is selected uniformly over the region on $[0, 1]^d$ that guarantees no overlap with the existing spheres. However, generating samples from this IS measure can be computationally challenging when $d \geq 2$. The grid-based IS introduced in Sections 4.4 and 4.5 overcomes this difficulty by imitating the reference IS; interestingly, it is more efficient than the reference IS.

In every algorithm presented in this paper, the running time complexity is calculated under the assumption that checking overlap of a newly generated sphere with an existing sphere is done in a sequential manner. That is, if there are n existing spheres, the expected running time complexity of the overlap check is proportional to n . However, if enough computing resources are available, the overlap check can be done in parallel so that its running time complexity is a constant. We omit the discussion of this parallel overlap check because it is easy to modify the results to accommodate the parallel case, and this does not change the key conclusions of the paper.

4.1. Naive acceptance–rejection algorithm

Algorithm 4.1 is a naive AR algorithm for generating perfect samples of the Gibbs hard-sphere model. The basic idea of the algorithm is standard [11], and its correctness is straightforward and hence omitted.

Let \mathcal{T}_{NAR} be the expected running time complexity of Algorithm 4.1, where the running time complexity denotes the number of elementary operations performed by the algorithm; every elementary operation takes at most a fixed amount of time. Note that the acceptance probability of each iteration is $\mathcal{P}(\lambda)$. Thus the expected total number of iterations of the algorithm is $1/\mathcal{P}(\lambda)$. Suppose $C_{\text{itr}}(\lambda)$ is the expected running time complexity of an iteration. Then

$$\mathcal{T}_{\text{NAR}} = \frac{C_{\text{itr}}(\lambda)}{\mathcal{P}(\lambda)}. \quad (4.1)$$

We now establish bounds on \mathcal{T}_{NAR} , and then establish its asymptotic behavior as $\lambda \nearrow \infty$ using Theorem 3. In each iteration of Algorithm 4.1, spheres are generated in a sequential order until we see an overlap or a configuration with N non-overlapping spheres. The key

Algorithm 4.1 Naive AR Method

```

1: repeat
2:   Generate  $N \sim \text{Poi}(\lambda)$ 
3:    $\mathbf{X} \leftarrow \emptyset$ 
4:   if  $N \neq 0$  then
5:      $i \leftarrow 0$ 
6:     repeat
7:        $i \leftarrow i + 1$ 
8:       Generate  $Y_i$  independently and uniformly distributed on  $[0, 1]^d$ 
9:       Generate a copy  $R_i$  of  $R$  independently of everything else
10:       $\mathbf{X} \leftarrow \mathbf{X} \cup \{(Y_i, R_i/\lambda^\eta)\}$ 
11:     until  $i = N$  or  $\mathbf{X} \notin \mathcal{A}$ 
12:   end if
13: until  $\mathbf{X} \in \mathcal{A}$ 
14: return  $\mathbf{X}$ 

```

to proving Proposition 4.1 is to establish that the expected number of spheres generated per iteration is $\Theta(\lambda^{\min\{\eta d, 2\}})$.

Proposition 4.1. *The expected running time complexity $C_{\text{itr}}(\lambda)$ of an iteration of the naive AR algorithm, Algorithm 4.1, satisfies*

$$C_{\text{itr}}(\lambda) = \Theta\left(\lambda^{\min\{\eta d, 2\}}\right). \tag{4.2}$$

Furthermore, the expected total running time \mathcal{T}_{NAR} satisfies

$$\mathcal{T}_{\text{NAR}} = \begin{cases} \Theta(\lambda^2) & \text{if } \eta d \geq 2, \\ \Theta\left(\lambda^{\eta d} \exp\left((\gamma m_1/2 + o(1))\lambda^{2-\eta d}\right)\right) & \text{if } 1 < \eta d < 2, \\ \Theta\left(\lambda^{\eta d} \exp(\delta\lambda)\right), \text{ for some } 0 < \delta \leq 1, & \text{if } \eta d = 1, \\ \Theta\left(\lambda^{\eta d} \exp\left((1 + o(1))\lambda\right)\right) & \text{if } 0 < \eta d < 1. \end{cases}$$

Remark 4.1. From (4.2) and Theorem 3.1, we see that for large values of λ and for $\eta d < 2$, \mathcal{T}_{NAR} is mainly governed by $\mathcal{P}(\lambda)$, which can be very small for large λ . This suggests that any rejection-based perfect sampling algorithm with a significant improvement in the acceptance probability will have a significantly improved running time complexity.

4.2. Importance-sampling-based acceptance–rejection algorithm

A sequence of tuples

$$\left\{ (D_{n,k}, \mu_{n,k}, \sigma_{n,k})_{k=1}^K \right\}_{n \in \mathbb{N}_0}$$

with some $K < \infty$ is called a *stable* IS sequence if for each $n \in \mathbb{N}_0$, $(D_{n,k})_{k=1}^K$ is a partition of \mathcal{G}_n , and $(\mu_{n,k})_{k=1}^K$ is a sequence of probability measures such that μ^0 is absolutely continuous

with respect to $\mu_{n,k}$ on $D_{n,k} \cap \mathcal{A}$, and the corresponding likelihood ratio

$$L_{n,k}(\mathbf{x}_n) := \frac{d\mu^0}{d\mu_{n,k}}(\mathbf{x}_n)$$

satisfies

$$L_{n,k}(\mathbf{x}_n) \leq \sigma_{n,k} \leq 1 \quad \text{if } \mathbf{x}_n \in D_{n,k} \cap \mathcal{A},$$

for $k = 1, \dots, K$. Under the stability condition, for every measurable subset $\mathcal{B} \subseteq \mathcal{G}$,

$$\begin{aligned} \mu(\mathcal{B}) &\propto \mathbb{P}_{\mu^0}(\mathbf{X} \in \mathcal{B} \cap \mathcal{A}) = \sum_{n \in \mathbb{N}_0} e^{-\lambda} \frac{\lambda^n}{n!} \left(\sum_{k=1}^K \mathbb{P}_{\mu^0}(\mathbf{X}_n \in D_{n,k} \cap \mathcal{B} \cap \mathcal{A}) \right) \\ &= \sum_{n \in \mathbb{N}_0} e^{-\lambda} \frac{\lambda^n}{n!} \left(\sum_{k=1}^K \mathbb{E}_{\mu_{n,k}} [I(\mathbf{X}_n \in D_{n,k} \cap \mathcal{B} \cap \mathcal{A}) L_{n,k}(\mathbf{X}_n)] \right) \\ &= \sum_{n \in \mathbb{N}_0} e^{-\lambda} \frac{\lambda^n \tilde{\sigma}(n)}{n!} \left(\sum_{k=1}^K \frac{\sigma_{n,k}}{\tilde{\sigma}(n)} \mathbb{E}_{\mu_{n,k}} \left[\frac{I(\mathbf{X}_n \in D_{n,k} \cap \mathcal{B} \cap \mathcal{A}) L_{n,k}(\mathbf{X}_n)}{\sigma_{n,k}} \right] \right) \\ &= \sum_{n \in \mathbb{N}_0} \frac{\lambda^n \tilde{\sigma}(n)}{n!} \left(\sum_{k=1}^K \frac{\sigma_{n,k}}{\tilde{\sigma}(n)} \mathbb{P}_{\mu_{n,k}}(J = 1, \mathbf{X}_n \in D_{n,k} \cap \mathcal{B} \cap \mathcal{A}) \right), \end{aligned} \tag{4.3}$$

where $\tilde{\sigma}(n) := \sum_{k=1}^K \sigma_{n,k}$, $U \sim \text{Unif}(0, 1)$, and

$$J \sim \text{Bern} \left(\frac{L_{n,k}(\mathbf{X}_n)}{\sigma_{n,k}} \right).$$

Let M be a non-negative integer-valued random variable with the probability mass function (PMF) defined by

$$\mathbb{P}(M = m) = \frac{1}{C_\lambda} \frac{\lambda^m \tilde{\sigma}(m)}{m!}, \quad m \in \mathbb{N}_0, \tag{4.4}$$

where

$$C_\lambda := \sum_{n=0}^\infty \frac{\lambda^n \tilde{\sigma}(n)}{n!}.$$

The PMF (4.4) is well defined because $\mathbb{E}[\tilde{\sigma}(N)]$ is finite under the stability condition. Now consider Algorithm 4.2.

Proposition 4.2. *Algorithm 4.2 generates a perfect sample of the Gibbs hard-sphere model. Furthermore, let $N \sim \text{Poi}(\lambda)$. Then the probability of accepting the configuration generated in an iteration of Algorithm 4.2 is given by*

$$P_{\text{acc}}(\lambda) = \frac{\mathcal{P}(\lambda)}{\mathbb{E}[\tilde{\sigma}(N)]}. \tag{4.5}$$

Algorithm 4.2 IS-based AR method

1: **repeat**

2: Generate a sample of M with PMF (4.4)

3: Generate J_1 with PMF $\mathbb{P}(J_1 = k) = \sigma_{M,k}/\tilde{\sigma}(M), k = 1, \dots, K$

4: Generate a realization \mathbf{X} of M points under the measure μ_{M,J_1}

5: Generate

$$J_2 \sim \text{Bern} \left(\frac{L_{M,J_1}(\mathbf{X})I(\mathbf{X} \in D_{M,J_1} \cap \mathcal{A})}{\sigma_{M,J_1}} \right)$$

6: **until** $J_2 = 1$

7: **return** \mathbf{X}

We omit the proof of Proposition 4.2 because the correctness easily follows from (4.3), and (4.5) holds from the observation that

$$P_{\text{acc}}(\lambda) = \frac{1}{C_\lambda} \sum_{n \in \mathbb{N}_0} \frac{\lambda^n \tilde{\sigma}(n)}{n!} \left(\sum_{k=1}^K \frac{\sigma_{n,k}}{\tilde{\sigma}(n)} \mathbb{E}^{\mu_{n,k}} \left[\frac{L_{n,k}(\mathbf{X}_n)}{\sigma_{n,k}}; \mathbf{X}_n \in D_{n,k} \cap \mathcal{A} \right] \right).$$

Note that the expected number of iterations of Algorithm 4.2 is $1/P_{\text{acc}}(\lambda)$. Corollary 4.2 is an important and trivial consequence of Proposition 4.2.

Corollary 4.1. *For all stable IS sequences*

$$\left\{ (D_{n,k}, \mu_{n,k}, \sigma_{n,k})_{k=1}^K \right\}_{n \in \mathbb{N}_0}$$

with the same $\mathbb{E}[\tilde{\sigma}(N)] = \sum_{k=1}^K \mathbb{E}[\sigma_{N,k}]$, the expected number of iterations of Algorithm 4.2 is the same.

Suppose that $\tilde{C}_{\text{itr}}(\lambda)$ is the expected running time complexity of an iteration of Algorithm 4.2. Then the expected total running time of the algorithm is given by

$$\mathcal{T}_{\text{ISAR}} = \frac{\tilde{C}_{\text{itr}}(\lambda)\mathbb{E}[\tilde{\sigma}(N)]}{\mathcal{P}(\lambda)}, \tag{4.6}$$

where $N \sim \text{Poi}(\lambda)$. Recall that the acceptance probability of the naive AR method is $\mathcal{P}(\lambda)$. It is reasonable to seek a valid stable IS sequence

$$\left\{ (D_{n,k}, \mu_{n,k}, \sigma_{n,k})_{k=1}^K \right\}_{n \in \mathbb{N}_0}$$

so that $\tilde{C}_{\text{itr}}(\lambda)\mathbb{E}[\tilde{\sigma}(N)]$ is smaller than $C_{\text{itr}}(\lambda)$. In Subsections 4.4 and 4.5, we present applications of Algorithm 4.2 where $\mathcal{T}_{\text{ISAR}}$ is indeed much smaller than \mathcal{T}_{NAR} .

Remark 4.2. (*Extension of IS-based AR to general Gibbs point processes.*) Suppose that μ is the distribution of a Gibbs point process that is absolutely continuous with respect to μ^0 , with the corresponding Radon–Nikodym derivative given by

$$\frac{d\mu}{d\mu^0}(\mathbf{x}) = \frac{\exp(-\beta V(\mathbf{x}))}{Z}, \quad \mathbf{x} \in \mathcal{X},$$




FIGURE 1. Illustration of the reference IS method for a Euclidean-hard-sphere model on $[0, 1]^2$ with spheres of fixed radius \bar{r}/λ^η . In (a) (respectively, (b)), the grey region represents the blocking area when generating the second circle (respectively, when generating the third circle).

where the constant $\beta \in \mathbb{R}$ is known as inverse temperature, V is called the non-negative potential function, and $Z = \mathbb{E}_{\mu^0} [\exp(-\beta V(\mathbf{X}))]$ is the normalizing constant. If the stability condition holds true when $I(\mathbf{x}_n \in \mathcal{A})$ is replaced by $\exp(-\beta V(\mathbf{x}_n))$, then Algorithm 4.2 can generate perfect samples from μ if in line 4.2 of the algorithm,

$$J_2 \sim \text{Bern}\left(\frac{L_{M,J_1}(\mathbf{X}) \exp(-\beta V(\mathbf{X})) I(\mathbf{X} \in D_{M,J_1})}{\sigma_{M,J_1}}\right).$$

To see that the hard-sphere model is a special case of such a Gibbs point process, take $\beta > 0$ and assume that $V(\mathbf{x}) = 0$ if \mathbf{x} is a non-overlapping configuration of spheres; otherwise, $V(\mathbf{x}) = \infty$.

4.3. Reference importance sampling

We now introduce an IS measure, called reference IS and denoted by $\tilde{\mu}_n$ for each n , such that $\{(\mathcal{G}_n, \tilde{\mu}_n, \sigma_n)\}_{n \in \mathbb{N}_0}$ is a stable IS sequence (with $K = 1$) that can be used in Algorithm 4.2 for generating perfect samples of the hard-sphere model for an appropriate choice of the sequence $\{\sigma_n : n \in \mathbb{N}_0\}$. Under $\tilde{\mu}_n$, we first generate an i.i.d. sequence R_1, \dots, R_n identical in distribution to R , and then generate n spheres sequentially as follows. Generate the center of the first sphere uniformly distributed on $[0, 1]^d$. Suppose that $i - 1$ spheres are already generated. For the i th sphere generation, a subset $\mathcal{B}_i \subseteq [0, 1]^d$ is called the *blocking region* if \mathcal{B}_i is the largest set such that if the center Y_i of the i th sphere fell in this region (that is, $Y_i \in \mathcal{B}_i$), the i th sphere would overlap with one of the existing $i - 1$ spheres. The center of the i th sphere is generated with uniform distribution over the non-blocking region $[0, 1]^d \setminus \mathcal{B}_i$. If, for some sphere $i \leq n$, the entire space is blocked (that is, $\mathcal{B}_i = [0, 1]^d$), we select the centers of spheres i, \dots, n arbitrarily. Figure 1 illustrates this for $d = 2$ and $n = 1, 2$. In conclusion, $\tilde{\mu}_n$ is the distribution of an output of Algorithm 4.3.

Observe that μ^0 is absolutely continuous with respect to $\tilde{\mu}_n$ on $\mathcal{G}_n \cap \mathcal{A}$, and the associated likelihood ratio satisfies

$$\tilde{L}_n(\mathbf{x}_n) = \frac{d\mu^0}{d\tilde{\mu}_n}(\mathbf{x}_n) = \prod_{i=1}^n (1 - B_i) \tag{4.7}$$

Algorithm 4.3 Reference IS

```

1: Input: The total number of spheres  $n$ 
2:  $\mathbf{X} \leftarrow \emptyset$ 
3: if  $n \neq 0$  then
4:    $\mathcal{B}_0 = \emptyset$  and  $i \leftarrow 0$ 
5:   repeat
6:      $i \leftarrow i + 1$ 
7:     Generate a copy  $R_i$  of  $R$  independently of everything else so far generated
8:     if  $\mathcal{B}_i = [0, 1]^d$  then
9:       Select the center  $Y_i$  of the  $i$ th sphere arbitrarily over  $[0, 1]^d$ 
10:    else
11:      Identify the non-blocking region  $\mathcal{B}_i^c$ 
12:      Generate  $Y_i$  uniformly distributed over  $\mathcal{B}_i^c$ 
13:    end if
14:     $\mathbf{X} \leftarrow \mathbf{X} \cup \{(Y_i, R_i/\lambda^\eta)\}$ 
15:  until  $i = n$ 
16: end if
17: Return  $\mathbf{X}$ 

```

for all $\mathbf{x}_n \in \mathcal{G}_n \cap \mathcal{A}$ and $n \in \mathbb{N}_0$, where B_i is the volume of \mathcal{B}_i and $\tilde{L}_0 = 1$. Note that $\tilde{L}_n(\mathbf{x}_n) = 0$ if and only if $\mathbf{x}_n \notin \mathcal{A}$, because for any $\mathbf{x}_n \notin \mathcal{A}$, there exists $i \leq n$ such that $B_i = 1$.

Observe that the blocking volume added by the i th sphere is at least $\gamma' (R_i/\lambda^\eta)^d$ when it does not overlap with any of the existing spheres. This is because, for the torus-hard-sphere model, the entire volume within an accepted sphere is added to the blocking volume, and for the Euclidean-hard-sphere model, at least $1/2^d$ of an accepted sphere is added to the blocking volume. Thus,

$$B_i \geq \frac{\gamma'}{\lambda^{\eta d}} \sum_{j=1}^{i-1} R_j^d \tag{4.8}$$

for every configuration $\mathbf{x}_{i-1} \in \mathcal{G}_{i-1} \cap \mathcal{A}$. In particular, if all the spheres are of the same size with a fixed radius \bar{r} ,

$$I(\mathbf{x}_n \in \mathcal{A}) \tilde{L}_n(\mathbf{x}_n) \leq \prod_{i=1}^n \left(1 - (i-1) \frac{\gamma'}{\lambda^{\eta d}} \bar{r}^d \right)^+ =: \delta_n \tag{4.9}$$

for all $n \in \mathbb{N}_0$ and $\mathbf{x}_n \in \mathcal{G}_n$, where $x^+ = \max(0, x)$ and $\delta_0 = 1$. Then the stability condition is satisfied with $K = 1$, $D_{n,1} = \mathcal{G}_n$, $\mu_{n,1} = \tilde{\mu}_n$, and $\sigma_{n,1} = \delta_n$ for $n \in \mathbb{N}_0$. Thus, Algorithm 4.2 generates perfect samples of the fixed-radius hard-sphere model, and from Proposition 4.2, the corresponding acceptance probability is

$$P_{\text{acc}}(\lambda) = \frac{\mathcal{P}(\lambda)}{\mathbb{E}[\tilde{\sigma}_N]} = \frac{\mathcal{P}(\lambda)}{\mathbb{E}[\delta_N]}.$$

Remark 4.3. In dimension $d = 1$, spheres become line segments and thus it is easy to generate samples from the IS measure $\tilde{\mu}_n$. However, for $d \geq 2$, generating samples under the reference

IS is difficult, because every time a new sphere is generated, we need to know the volume of the blocking region created by the existing spheres, and then we need to generate a point uniformly on this non-blocking region; see line 11 in Algorithm 4.3. One possible way to implement the reference IS is by combining a well-known method called *power tessellation* and a simple rejection method in two steps: (i) using the power tessellation, we can compute the blocking volumes exactly (see e.g. [4, 33]); (ii) we then use a simple AR method, repeatedly generating a point independently and uniformly on $[0, 1]^d$ until it falls within the non-blocking region.

Unfortunately, implementing the power tessellation method is computationally prohibitive. Besides, even if we have an efficient implementation of the power tessellation method, the above simple rejection step can be expensive when the non-blocking region is small. Fortunately, we can overcome both these difficulties by using a simple grid on $[0, 1]^d$. From (4.6), it is evident that if there are two IS methods with the same $\mathbb{E}[\tilde{\sigma}_N]$, it is computationally preferable to use the method that has smaller per-iteration expected running time, $\tilde{C}_{\text{itr}}(\lambda)$. In Subsection 4.4, we introduce a hypercubic-grid-based IS method that continues to generate perfect samples while the blocking regions are closely approximated by grid cells. With a careful choice of the cell-edge length, we make sure that the inequality (4.9) holds for the grid IS as well (and thus $\mathbb{E}[\tilde{\sigma}_N]$ is same as that of the reference IS). As a consequence of Corollary 4.1, the expected number of iterations of Algorithm 4.2 is the same as that of the reference IS method. However, the grid method is easy to implement and has a much smaller expected iteration cost $\tilde{C}_{\text{itr}}(\lambda)$ than the reference IS. The choice of the hypercubic grid is just an option that simplifies the implementation; however, the method can be implemented using other kinds of grids. In the two-dimensional case, for example, it is possible to use a hexagonal grid to implement the IS method.

4.4. Grid-based importance sampling for fixed-radius case

Consider the hard-sphere model with a fixed radius \bar{r}/λ^η . The process of generating n spheres under the following grid-based IS measure $\hat{\mu}_n$ starts with a partitioning of the underlying space $[0, 1]^d$ into a hypercubic grid with a cell-edge length $\varepsilon > 0$ such that $1/\varepsilon$ is an integer. The centers of the spheres are generated sequentially, as follows. Suppose that $i - 1$ spheres with centers Y_1, \dots, Y_{i-1} have already been generated. At the time of generation of the i th sphere, a cell C in the grid is labeled as *fully-blocked* if the cell is completely inside a sphere with radius $2\bar{r}/\lambda^\eta$ centered at an existing point, that is, $C \subseteq S(Y_j, 2\bar{r}/\lambda^\eta)$ for some $j \leq i - 1$; otherwise, the cell is labeled as *non-fully-blocked*. A non-fully-blocked cell C is called *partially-blocked* if $C \cap S(Y_j, 2\bar{r}/\lambda^\eta) \neq \emptyset$ for some $j \leq i - 1$; otherwise, it is called *non-blocked*. The center Y_i of the i th sphere is selected uniformly over the non-fully-blocked cells, because selecting Y_i over a fully-blocked cell will certainly result in the i th sphere overlapping with an existing sphere. We then check for overlap only if Y_i is generated over a partially-blocked cell, because overlap is not possible if Y_i is generated over a non-blocked cell. If either there is an overlap or all the cells are fully-blocked by the existing spheres, the centers Y_i, \dots, Y_n of the remaining spheres are selected arbitrarily (such a selection results in an overlapping configuration). Otherwise, for the generation of the next sphere, $i + 1$, we repeat the same procedure by relabeling the non-fully-blocked cells, considering spheres $1, \dots, i$ as the existing spheres. Note that at the beginning of each iteration all the cells are labeled as non-blocked. Also note that since all the spheres have the same radius, for the relabeling of the cells we only need to focus on the cells that might interact with the last sphere generated. See Figure 2 for an illustration of this sequential procedure.




FIGURE 2. A realization with five circles on the unit square $[0, 1]^2$ generated using the grid-based IS method for a Euclidean-hard-sphere model with a fixed radius (smaller circles). The grid size is 50×50 and the radius is 0.1. The bigger circle around each point is the actual region blocked by the circle. For the sixth circle generation, grey cells are fully-blocked, hatched cells are partially-blocked, and white cells are non-blocked.

Suppose that $\widehat{\mu}_n$ is the probability measure under which n spheres are generated by the above procedure. Then μ^0 is absolutely continuous with respect to $\widehat{\mu}_n$ on $\mathcal{G}_n \cap \mathcal{A}$ and the corresponding likelihood ratio is

$$\widehat{L}_n(\mathbf{x}_n) := \frac{d\mu^0}{d\widehat{\mu}_n}(\mathbf{x}_n) = \prod_{i=1}^n (1 - \widehat{B}_i), \quad \mathbf{x}_n \in \mathcal{G}_n \cap \mathcal{A},$$

where \widehat{B}_i is the volume of fully-blocking cells for the i th sphere generation; that is, \widehat{B}_i equal to the product of the number of fully-blocked cells and ε^d . To apply Algorithm 4.2 for the fixed-radius hard-sphere model, take $K = 1$, and for each $n \in \mathbb{N}_0$, take $D_{n,1} = \mathcal{G}_n$, $\mu_{n,1} = \widehat{\mu}_n$, and $\sigma_{n,1} = \delta_n$. Thus, $\widetilde{\sigma}(n) = \delta_n$ and $L_{n,1}(\mathbf{x}_n) = \widehat{L}_n(\mathbf{x}_n)$ for all $\mathbf{x}_n \in \mathcal{G}_n \cap \mathcal{A}$.

Selection of the cell-edge length ε : Observe that the longest diagonal length of a cell is $\sqrt{d} \varepsilon$. Since we focus only on the non-overlapping configurations, in the implementation, we generate a sphere only if all the existing spheres are non-overlapping. Suppose that the cell-edge length ε is selected so that $\sqrt{d} \varepsilon \leq \overline{r}/\lambda^\eta$. Then for the i th sphere generation, every cell that has non-empty intersection with $S(Y_j, \overline{r}/\lambda^\eta)$, for any $j = 1, \dots, i - 1$, has to be fully-blocked, because such a cell is a subset of $S(Y_j, 2\overline{r}/\lambda^\eta)$. Thus, the non-overlapping condition of the existing spheres imply that

$$\cup_{j=1}^{i-1} S(Y_j, \overline{r}/\lambda^\eta)$$

is a subset of the union of the fully-blocked cells, and hence

$$\widehat{B}_i \geq \frac{\gamma'(i-1)\bar{r}^d}{\lambda^{\eta d}}.$$

Thus, for $n \geq 1$,

$$I(\mathbf{x}_n \in \mathcal{A})\widehat{L}_n(\mathbf{x}_n) \leq \delta_n = \prod_{i=1}^n \left(1 - (i-1)\frac{\gamma'\bar{r}^d}{\lambda^{\eta d}}\right)^+, \quad \mathbf{x}_n \in \mathcal{G}_n. \tag{4.10}$$

This upper bound is same as the one we obtained in the case of the reference IS; see the inequality (4.9). Since the acceptance probability $P_{\text{acc}}(\lambda) = \mathcal{P}(\lambda)/\mathbb{E}[\delta_N]$ is the same for both the grid IS and the reference IS methods, we need to choose the cell-edge length $\varepsilon \leq \bar{r}/\lambda^\eta$ so that the expected per-iteration running time $\widetilde{C}_{\text{itr}}(\lambda)$ is minimum. It is easy to see that the higher the value of ε , the smaller $\widetilde{C}_{\text{itr}}(\lambda)$, for the following reasons:

1. Labeling of the cells is faster if they are bigger in size.
2. Increasing the cell size increases the chances of overlap of the new sphere with the existing spheres, and hence on average each iteration generates fewer spheres.

In conclusion, we choose $\varepsilon = 1/\lfloor \lambda^\eta/\bar{r} \rfloor$ for the implementation of the grid IS method.

To reduce the per-iteration complexity of the algorithm, we make some changes to Steps 4 and 5 in Algorithm 4.2. Observe that a realization \mathbf{X}_n generated under $\widehat{\mu}_n$ is accepted only if $\mathbf{X}_n \in \mathcal{A}$ and $J = 1$, where $J \sim \text{Bern}(\widehat{L}_n(\mathbf{X}_n)/\delta_n)$. In the implementation, we generate an i.i.d. sequence $U_1, \dots, U_n \sim \text{Unif}(0, 1)$ independent of everything else so far generated, and take

$$J_i = I\left(U_i \leq \frac{1 - \widehat{B}_i}{(1 - (i-1)\gamma'\bar{r}^d\lambda^{-\eta d})}\right)$$

for $i \leq n$. Since J and the product $\prod_{i=1}^n J_i$ are Bernoulli random variables with the same success probability $L(\mathbf{X}_n)/\delta_n$, to reduce the per-iteration cost, we generate the i th sphere only if $J_i = 1$ and the existing spheres do not overlap with each other.

Algorithm 4.4 implements the grid-based IS for a given n with the above-mentioned enhancements. Algorithm 4.2 is restated as Algorithm 4.5.

Algorithm 4.4 Grid-based IS for fixed radius

- 1: **Input:** The total number of spheres $n \geq 1$ and a grid on $[0, 1]^d$
 - 2: **Output:** $(\mathbf{X}, \text{Status}) \in \mathcal{G} \times \{\text{True}, \text{False}\}$. Where $\text{Status} = \text{True}$ if $\mathbf{X} \in \mathcal{A}$ and $\text{Status} = \text{False}$ otherwise
 - 3: Label every cell as non-blocked
 - 4: $\mathbf{X} \leftarrow \emptyset, i \leftarrow 0$ and $\widehat{B} \leftarrow 0$
 - 5: **repeat**
 - 6: $i \leftarrow i + 1$
 - 7: Generate $U \sim \text{Unif}(0, 1)$
 - 8: **if** $U > \frac{1 - \widehat{B}}{1 - (i-1)\gamma'\bar{r}^d\lambda^{-\eta d}}$ **then**
 - 9: **return** $(\mathbf{X}, \text{False})$
-

```

10:  else
11:      Generate  $Y_i$  uniformly distributed over the non-fully-blocked cells
          (and independently of everything else so far generated)
12:      if  $Y_i$  is on a partially-blocked cell and there is an overlap then
13:          return  $(\mathbf{X}, \text{False})$ 
14:      end if
15:      Update the cell labels
16:      Compute the volume  $\widehat{B}$  of the fully-blocked cells
17:  end if
18:   $\mathbf{X} \leftarrow \mathbf{X} \cup \{(Y_i, \bar{r}/\lambda^\eta)\}$ 
19: until  $i = n$ 
20: return  $(\mathbf{X}, \text{True})$ 

```

Algorithm 4.5 Perfect sampling for hard-sphere model using grid-based IS

```

1: Partition  $[0, 1]^d$  into a hypercube grid with cell-edge length  $\varepsilon = 1/\lfloor \lambda^\eta/\bar{r} \rfloor$ 
2: repeat
3:   Generate a sample of  $M$  with PMF (4.4)
4:   if  $M = 0$  then
5:      $(\mathbf{X}, \text{Status}) \leftarrow (\emptyset, \text{True})$ 
6:   else
7:     Obtain an output  $(\mathbf{X}, \text{Status})$  from Algorithm 4.4 with  $M$  and the grid as input
8:   end if
9: until Status = True
10: return  $\mathbf{X}$ 

```

Remark 4.4. (*The PMF of M .*) Note that, for the current setup, the PMF of M , given by (4.4), becomes

$$\mathbb{P}(M = m) = \frac{1}{C_\lambda} \frac{\lambda^m \delta_m}{m!},$$

$m \in \mathbb{N}_0$, where the normalizing constants are

$$C_\lambda = \sum_{n \in \mathbb{N}_0} \frac{\lambda^n \delta_n}{n!}.$$

The support of the PMF is finite because $\delta_m = 0$ for all $m \geq \lambda^{\eta d}/(\gamma' \bar{r}^d) + 1$. To increase the performance of the algorithm, we can further truncate the support of the PMF. Using the maximum packing density, we can obtain an integer m_{\max} such that $\mathbf{X} \notin \mathcal{A}$ for all $m \geq m_{\max}$ and configurations \mathbf{X} with $|\mathbf{X}| = m$. In that case, we can take

$$\mathbb{P}(M = m) = \frac{1}{C_\lambda} \frac{\lambda^m \delta_m}{m!}, \quad 0 \leq m \leq m_{\max},$$

with

$$C_\lambda = \sum_{n=0}^{m_{\max}} \frac{\lambda^n \delta_n}{n!}.$$

For example, refer to [29] for finding maximum packing densities for $d = 2$ and $d = 3$.

We now focus on the expected running time analysis of Algorithm 4.5. By Proposition 4.2, the acceptance probability $P_{\text{acc}}(\lambda)$ of Algorithm 4.5 is $\mathcal{P}(\lambda)/\mathbb{E}[\tilde{\sigma}(N)] = \mathcal{P}(\lambda)/\mathbb{E}[\delta_N]$. A proof of Proposition 4.3 is given in Section A.4.

Proposition 4.3. *For the fixed-radius hard-sphere model, there exists a constant $c > 0$ such that*

$$\mathcal{T}_{\text{ISAR}} \leq c \mathbb{E}[\delta_N] \frac{\lambda^{\min\{\eta d, 1\}}}{\mathcal{P}(\lambda)}, \tag{4.11}$$

where $N \sim \text{Poi}(\lambda)$. Furthermore,

$$\begin{aligned} \limsup_{\lambda \nearrow \infty} \left[\frac{1}{\lambda^{2-\eta d}} \log \mathbb{E}[\delta_N] \right] &\leq -\frac{\gamma' \bar{r}^d}{2} \text{ if } \eta d > 1, \text{ and} \\ \limsup_{\lambda \nearrow \infty} \left[\frac{1}{\lambda} \log \mathbb{E}[\delta_N] \right] &\leq -b \text{ if } 0 < \eta d \leq 1, \text{ for some constant } b > 0. \end{aligned}$$

The following result is a trivial consequence of Propositions 4.1 and 4.3.

Corollary 4.2. *For the fixed-radius hard-sphere model, if $\eta d \geq 2$, both $\mathcal{T}_{\text{ISAR}}$ and \mathcal{T}_{NAR} are of the same order; and if $0 < \eta d < 2$, there exists a constant $c > 0$ such that $\mathcal{T}_{\text{ISAR}} \leq c \mathbb{E}[\delta_N] \mathcal{T}_{\text{NAR}}$.*

Remark 4.5. (Better choice of δ_n for the Euclidean-hard-sphere model) If the spheres are Euclidean, further improvements in the choice of δ_n can be obtained by accounting for boundary effects. For instance, for $d = 2$, the four corners of $[0, 1]^2$ are covered by at most four circles, each of which contributing a blocking area of at least $\gamma' \bar{r}^2 / \lambda^{2\eta} = \pi \bar{r}^2 / 4\lambda^{2\eta}$, while each of the remaining circles contributing a blocking area of at least $2\gamma' \bar{r}^2 / \lambda^{2\eta} = \pi \bar{r}^2 / 2\lambda^{2\eta}$. Let $b_0 = 0$,

$$b_i = (i - 1) \frac{\pi \bar{r}^2}{4\lambda^{2\eta}}$$

for $1 \leq i \leq 5$, and

$$b_i = \frac{\pi \bar{r}^2}{\lambda^{2\eta}} + (i - 4) \frac{\pi \bar{r}^2}{2\lambda^{2\eta}}$$

for $i \geq 6$. Then, for this particular scenario, a better choice of δ_n in (4.9) (as well as in (4.10)) is $\delta_n = \prod_{i=1}^n (1 - b_n)^+$, $n \in \mathbb{N}_0$.

4.5. Random-radii case

We now consider another application of Algorithm 4.2 for the hard-sphere model when under the marked PPP the radii of the spheres are i.i.d. For the fixed-radius case presented in Section 4.4, the proposed IS method ensured a uniform bound δ_n on the likelihood ratio over \mathcal{G}_n for every $n \in \mathbb{N}_0$, as shown in (4.10). Such upper bounds are possible for a random-radii hard-sphere model if the radii are bounded below by a positive constant. Furthermore, a similar analysis can be established when the spheres are replaced with i.i.d. convex shapes such that each shape occupies a minimum positive volume. However, when the radii are not bounded from below almost surely, the associated blocking volumes can be arbitrarily small. We address this issue by partitioning \mathcal{G}_n into two sets $D_{n,1}$ and $D_{n,2}$ for each n so that the IS on

$D_{n,1}$ is a grid-based IS method that is similar to Algorithm 4.4, and the IS on $D_{n,2}$ is obtained by exponentially twisting the distribution of R^d to put high probability mass on configurations with lower-volume spheres.

We first introduce the exponential twisting of the distribution, say G , of R^d . Recall that R is assumed to be a bounded non-negative random variable. Without loss of generality further assume that $\alpha := \mathbb{E}[R^d] > 0$. Thus the logarithmic moment generating function of R^d , defined by $\Lambda(\theta) := \log(\mathbb{E}[\exp(\theta R^d)])$, is finite for every $\theta \in \mathbb{R}$. Furthermore, the derivative

$$\Lambda'(\theta) = \frac{d\Lambda(\theta)}{d\theta} = \frac{\mathbb{E}[R^d \exp(\theta R^d)]}{\mathbb{E}[\exp(\theta R^d)]}$$

is finite and positive for all $\theta \in \mathbb{R}$, and in particular, $\Lambda'(0) = \alpha$. In fact, using the results in Chapter 2 of [10], it can be seen that $\Lambda(\theta)$ is strictly convex. As a consequence, $\Lambda'(\theta)$ is strictly increasing and hence

$$\alpha_{\min} := \lim_{\theta \rightarrow -\infty} \Lambda'(\theta) < \alpha.$$

Let $\hat{\theta}$ be such that $\Lambda'(\hat{\theta}) = \varrho$ for some $\varrho \in (\alpha_{\min}, \alpha)$. Therefore, $\hat{\theta} < 0$. Now consider the distribution \tilde{G} obtained by exponentially twisting G by the amount $\hat{\theta}$, that is, $d\tilde{G}(t) = \exp(\hat{\theta}t - \Lambda(\hat{\theta})) dG(t)$. Fix a constant $a \in (0, 1)$, and for each integer $n \geq 1$, define

$$H_n := \left\{ (t_1, t_2, \dots, t_{\lceil na \rceil}) \in \mathbb{R}_+^{\lceil na \rceil} : \frac{1}{\lceil na \rceil} \sum_{i=1}^{\lceil na \rceil} t_i < \varrho \right\}.$$

We later see that $a = 1/2$ is a good choice for increasing performance of the algorithm. Let $\Lambda^*(\cdot)$ be the Legendre–Fenchel transform of Λ , that is, $\Lambda^*(t) = \sup_{\theta \in \mathbb{R}} \{\theta t - \Lambda(\theta)\}$. This corresponds to the large deviations rate function associated with the empirical average of i.i.d. samples from G . From the definition of $\hat{\theta}$ and the fact that $\Lambda(\theta)$ is strictly convex, $\Lambda^*(\varrho) = \hat{\theta}\varrho - \Lambda(\hat{\theta}) > 0$. Since $\hat{\theta} < 0$, for all $(t_1, t_2, \dots, t_{\lceil na \rceil}) \in H_n$,

$$\exp\left(\hat{\theta} \sum_{i=1}^{\lceil na \rceil} t_i - \lceil na \rceil \Lambda(\hat{\theta})\right) = \exp\left(\hat{\theta} \sum_{i=1}^{\lceil na \rceil} (t_i - \varrho) + \lceil na \rceil \Lambda^*(\varrho)\right) \geq \exp(\lceil na \rceil \Lambda^*(\varrho)),$$

and thus,

$$\prod_{i=1}^{\lceil na \rceil} \frac{dG}{d\tilde{G}}(t_i) \leq \exp(-\lceil na \rceil \Lambda^*(\varrho)) \leq \exp(-na \Lambda^*(\varrho)). \tag{4.12}$$

Recall the definition of the distribution μ of the hard-sphere model given by (2.1). To apply Algorithm 4.2, select $K = 2$ and define

$$D_{n,1} = \left\{ \mathbf{x}_n = \{(x_1, r_1/\lambda^{\eta d}), \dots, (x_n, r_n/\lambda^{\eta d})\} \in \mathcal{G}_n : (r_1^d, \dots, r_{\lceil na \rceil}^d) \in H_n^c \right\}$$

and $D_{n,2} = \mathcal{G}_n \setminus D_{n,1}$, for each n , where H_n^c is the complement of H_n within $[0, \bar{r}]^{\lceil na \rceil}$. To apply Algorithm 4.5, we are now left with identifying the IS measures $\mu_{n,1}$ and $\mu_{n,2}$ and the corresponding bounds $\sigma_{n,1}$ and $\sigma_{n,2}$ for each $n \in \mathbb{N}_0$.

The measure $\mu_{n,1}$ on $D_{n,1}$ is again a grid-based IS method similar to the grid method introduced for the fixed-radius case in Subsection 4.4. First, i.i.d. copies $\{R_1, \dots, R_n\}$ of R are

generated. Then we construct a new grid and label each cell every time a new sphere is generated, as follows. For the generation of the i th, sphere with radius R_i/λ^η , we take the cell-edge length $\varepsilon = 1/\lceil \lambda^\eta/R_i \rceil$. A cell C in the grid is labeled as fully-blocked if $C \subseteq S(X_j, (R_j + R_i)/\lambda^\eta)$ for an existing sphere $j \leq i - 1$ with the center X_j and the radius R_j/λ^η ; otherwise, the cell is labeled as non-fully-blocked. A non-fully-blocked cell C is called partially-blocked if $C \cap S(X_j, (R_j + R_i)/\lambda^\eta) \neq \emptyset$ for some $j \leq i - 1$; otherwise, it is called non-blocked. Then the next center X_i is generated uniformly over the non-fully-blocked cells. Just as in the case of fixed radius, X_1 is generated uniformly over $[0, 1]^d$, and we check the possibility of the overlap of i th sphere with an existing sphere only if X_i falls in a partially-blocked cell.

The measure μ^0 is absolutely continuous with respect to $\mu_{n,1}$ on $D_{n,1} \cap \mathcal{A}$, and the associated likelihood ratio $L_{n,1}$ is given by

$$L_{n,1}(\mathbf{x}_n) = \prod_{i=1}^n (1 - B_i), \quad \mathbf{x}_n \in D_{n,1} \cap \mathcal{A},$$

where \widehat{B}_i is the volume of all the fully-blocked cells for the generation of the i th sphere. By (4.8) and the fact that the cell-edge length is $1/\lceil \lambda^\eta/R_i \rceil$, we have

$$B_i \geq \min \left(1, \frac{\gamma' \lceil na \rceil}{\lambda^{\eta d}} \varrho \right)$$

on $D_{n,1} \cap \mathcal{A}$ for all $i \geq \lceil na \rceil + 1$, because

$$\frac{1}{\lceil na \rceil} \sum_{j=1}^{\lceil na \rceil} R_j^d \geq \varrho$$

over the set H_n^c . Consequently,

$$I(\mathbf{x}_n \in \mathcal{A})L_{n,1}(\mathbf{x}_n) \leq \left[\left(1 - \frac{\gamma' \lceil na \rceil}{\lambda^{\eta d}} \varrho \right)^+ \right]^{n(1-a)} =: \sigma_{n,1}, \quad \mathbf{x}_n \in D_{n,1}.$$

The measure $\mu_{n,2}$ is induced by the following procedure. Generate i.i.d. samples $R_1^d, \dots, R_{\lceil na \rceil}^d$ from \widetilde{G} , and independently of this, generate i.i.d. samples $R_{\lceil na \rceil+1}^d, \dots, R_n^d$ from G . For $i = 1, \dots, n$, the radius of the i th sphere is R_i/λ^d , and the center is generated uniformly distributed over the non-blocking region created by the existing $i - 1$ spheres. Since $R_1^d, \dots, R_{\lceil na \rceil}^d$ are sampled from \widetilde{G} , by (4.12),

$$I(\mathbf{x}_n \in \mathcal{A})L_{n,2}(\mathbf{x}_n) = \prod_{i=1}^{\lceil na \rceil} \frac{dG}{d\widetilde{G}}(r_i^d) \leq \exp(-na \Lambda^*(\varrho)) =: \sigma_{n,2} \quad \text{for all } \mathbf{x}_n \in D_{n,2}.$$

In summary,

$$\left\{ (D_{n,k}, \mu_{n,k}, \sigma_{n,k})_{k=1}^2 \right\}_{n \in \mathbb{N}_0}$$

is a stable IS sequence, and hence Algorithm 4.2 generates perfect samples from μ . However, to reduce the per-iteration complexity (as in the fixed-radius case), we make some modifications to the algorithm. Algorithm 4.6 is similar to Algorithm 4.4, and Algorithm 4.2 is restated as Algorithm 4.7.

Algorithm 4.6 Grid-based IS for random-radii case

1: **Input:** The total number of spheres $n \geq 1$
 2: **Output:** $(\mathbf{X}, \text{Status}) \in \mathcal{G} \times \{\text{True}, \text{False}\}$. Where $\text{Status} = \text{True}$ if $\mathbf{X} \in \mathcal{A}$ and $\text{Status} = \text{False}$ otherwise
 3: $i \leftarrow 0$
 4: $\mathbf{X} \leftarrow \emptyset$
 5: **repeat**
 6: $i \leftarrow i + 1$
 7: Generate a copy R_i of R independently of everything else so far generated
 8: Construct a grid on $[0, 1]^d$ with the cell-edge length $\varepsilon = 1/\lceil \lambda^\eta / R_i \rceil$
 9: Identify the label of each cell in the new grid
 10: Compute the volume \widehat{B} of the fully-blocked cells and generate $U \sim \text{Unif}(0, 1)$
 11: **if**

$$U > \frac{1 - \widehat{B}}{(1 - \gamma' \lceil na \rceil \varrho \lambda^{-nd})^{1-a}}$$

then
 12: **return** $(\mathbf{X}, \text{False})$
 13: **else**
 14: Generate Y_i uniformly distributed over the non-fully-blocked cells
 (and independently of everything else so far generated)
 15: **if** Y_i is on a partially-blocked cell and there is an overlap **then**
 16: **return** $(\mathbf{X}, \text{False})$
 17: **end if**
 18: **end if**
 19: $\mathbf{X} \leftarrow \mathbf{X} \cup \{(Y_i, R_i/\lambda^\eta)\}$
 20: **until** $i = n$
 21: **return** $(\mathbf{X}, \text{True})$

We now focus on the running time complexity of Algorithm 4.7. Notice that $\tilde{\sigma}(n) = \sigma_{n,1} + \sigma_{n,2}$ for each $n \in \mathbb{N}_0$. By Proposition 4.2, $P_{\text{acc}}(\lambda) = \mathcal{P}(\lambda)/\mathbb{E}[\tilde{\sigma}(N)]$ with $N \sim \text{Poi}(\lambda)$. Observe that

$$\sigma_{n,1} \leq \exp\left(-\frac{\gamma' n^2 a(1-a)}{\lambda^{\eta d}} \varrho\right).$$

The proof of Proposition 4.3 can be extended to the current scenario to show that

$$\limsup_{\lambda \nearrow \infty} \left[\frac{1}{\lambda^{2-\eta d}} \log \mathbb{E}[\sigma_{N,1}] \right] \leq -\gamma' a(1-a)\varrho \quad \text{if } \eta d > 1, \text{ and}$$

$$\limsup_{\lambda \nearrow \infty} \left[\frac{1}{\lambda} \log \mathbb{E}[\sigma_{N,1}] \right] \leq -b \quad \text{if } 0 < \eta d \leq 1, \text{ for some constant } b > 0.$$

It is now clear that a good choice for a is $1/2$, because it maximizes $a(1-a)$. Furthermore, using the moment generating function of Poisson random variables, we have

$$\mathbb{E}[\sigma_{N,2}] \leq \exp\left(-\lambda \left(1 - e^{-\Lambda^*(\varrho/2)}\right)\right).$$

Algorithm 4.7 Perfect sampling for hard-sphere model with random radii

```

1: repeat
2:   Generate a sample of  $M$  with PMF (4.4)
3:   Generate  $J$  with PMF  $\mathbb{P}(J = k) = \sigma_{M,k}/\tilde{\sigma}(M)$ ,  $k = 1, 2$ 
4:   if  $J = 1$  then
5:     Obtain an output  $(\mathbf{X}, \text{Status})$  of Algorithm 4.6 with  $M$  as input
6:   else
7:     Generate  $\mathbf{X}$  under  $\mu_{M,2}$ 
8:     if  $\text{Bern}\left(\frac{L_{M,2}(\mathbf{X})I(\mathbf{X} \in D_{M,2} \cap \mathcal{A})}{\sigma_{M,2}}\right) = 0$  then
9:       Status  $\leftarrow$  False
10:    end if
11:  end if
12: until Status = True
13: return  $\mathbf{X}$ 

```

Recall that $\mathcal{T}_{\text{ISAR}} \leq \mathbb{E}[\tilde{\sigma}(N)] \tilde{C}_{\text{itr}}(\lambda)/\mathcal{P}(\lambda)$. The per-iteration complexity $\tilde{C}_{\text{itr}}(\lambda)$ is mainly determined by the relabeling of cells in the new grid for each sphere generation. The grid size for the i th sphere generation is of order $\lambda^{\eta d}/R_i^d$, and the total number of spheres generated in each iteration is at most of order $\lambda^{\min\{\eta d, 1\}}$. Therefore, for $\eta d \leq 2$, we can show that $\tilde{C}_{\text{itr}}(\lambda)$ is of order $\lambda^{\min\{\eta d, 1\}}\mathbb{E}[1/R^d]$.

Remark 4.6. If ϱ is selected to equal

$$\text{argmin}_{\varrho \in (\alpha_{\min}, \alpha)} (\sigma_{n,1} + \sigma_{n,2})$$

for each $n = 1, 2, \dots$, then $\mathbb{E}[\tilde{\sigma}(N)]$ is minimum. Note that $\sigma_{n,1}$ decreases and $\sigma_{n,2}$ increases as functions of ϱ . The above decompositions were chosen to illustrate ideas simply. More complex decompositions are easily created for further performance improvement. For instance, we could have defined H_n above as

$$H_n^c := \left\{ (r_1, \dots, r_n) \in \mathbb{R}_+^n : \frac{1}{m} \sum_{i=1}^m r_i \geq \varrho_m, \forall m \leq n \right\},$$

and then arrived at appropriate $\{\varrho_m\}_{m \leq n}$ and appropriate changes of measures for configurations in H_n and H_n^c . While this should lead to substantial performance improvement, it also significantly complicates the analysis.

5. Dominated coupling-from-the-past methods

In this section, we review some of the well-known dominated CFTP algorithms for the hard-sphere models. We refer to [27] for a general description of dominated CFTP for Gibbs point processes (this method was first proposed for area-interaction processes by Kendall [25]).

Let $\mathbf{D} = \{\mathbf{D}(t) : t \in \mathbb{R}\}$ be the so-called dominating birth-and-death process on $[0, 1]^d$, with births arriving as a Poisson process with rate λ , where each birth is a uniformly and

independently generated marked point on $[0, 1]^d$ denoting the center of a sphere, with the mark being its radius. Each birth is alive for an independent random time exponentially distributed with mean one. It is well known that the steady-state distribution of \mathbf{D} is μ^0 . Furthermore, it is easy to generate the dominating process \mathbf{D} both forward and backward in time so that $\mathbf{D}(t) \sim \mu^0$ for all t . To see this, let $\dots < t_{-2} < t_{-1} < 0 < t_1 < t_2 < \dots$ be the event instants of the process \mathbf{D} , where an event can be either a birth or a death. Assume that with each birth there is an additional mean-one exponentially distributed independent mark to determine its lifetime. Since the births are arriving as a Poisson process, the interarrival times are exponential with mean $1/\lambda$. Generate $\mathbf{D}(0) \sim \mu^0$, determine the next event instant t_1 , and take $\mathbf{D}(t) = \mathbf{D}(0)$ for $0 \leq t < t_1$. If the next event is a birth, generate a new independent (marked) point; otherwise, remove the existing point with the smallest lifetime. Continue the same procedure starting with $\mathbf{D}(t_1)$ to generate the process over $[t_1, t_2)$, and so on. For generating the dominating process \mathbf{D} backward in time, observe that \mathbf{D} is time-reversible, and hence we can generate $\{\mathbf{D}(t) : -T \leq t \leq 0\}$ for any finite $T > 0$ just by generating an independent copy $\{\tilde{\mathbf{D}}(t) : 0 \leq t \leq T\}$ of the dominating process $\{\mathbf{D}(t) : 0 \leq t \leq T\}$ and taking $\mathbf{D}(-t) = \tilde{\mathbf{D}}(t)$ for $0 \leq t \leq T$.

Since the distribution μ of the hard-sphere model is absolutely continuous with respect to μ^0 , using *coupling*, it is possible to construct a spatial birth-and-death process $\mathbf{Z} = \{\mathbf{Z}(t) : t \in \mathbb{R}\}$, called the *interaction* process, such that $\mathbf{Z}(t) \subseteq \mathbf{D}(t)$ and $\mathbf{Z}(t) \sim \mu$ for all $t \in \mathbb{R}$; see [27]. Each iteration of any dominated CFTP method essentially involves the following two steps:

1. Fix $n > 0$ and construct $\{\mathbf{D}(t) : t_{-n} \leq t \leq 0\}$ backward in time starting at time zero with $\mathbf{D}(0) \sim \mu^0$.
2. Then, as detailed in Sections 5.1–5.3, use *thinning* on the dominating process $\{\mathbf{D}(t) : t_{-n} \leq t \leq 0\}$ to obtain an upper bounding process $\{\mathbf{U}_n(t) : t \geq t_{-n}\}$ with $\mathbf{U}_n(t_{-n}) = \mathbf{D}(t_{-n})$ and a lower bounding process $\{\mathbf{L}_n(t_{-n}) : t \geq t_{-n}\}$ with $\mathbf{L}_n(t_{-n}) = \emptyset$ forward in time, such that for $t \geq t_{-n}$, $\mathbf{L}_n(t) \subseteq \mathbf{Z}(t) \subseteq \mathbf{U}_n(t) \subseteq \mathbf{D}(t)$ and $\mathbf{L}_m(t) \subseteq \mathbf{L}_n(t) \subseteq \mathbf{U}_n(t) \subseteq \mathbf{U}_m(t)$ for $m \leq n$.

If \mathbf{U}_n and \mathbf{L}_n *coalesce* at time 0, that is, $\mathbf{U}_n(0) = \mathbf{L}_n(0)$, then $\mathbf{U}_n(0)$ is a perfect sample from the target distribution μ . If there is no coalescence, then repeat the steps by increasing n and extending the dominating process $\{\mathbf{D}(t) : -t_{-n} \leq t \leq 0\}$ further backward to time t_{-n} and repeat the same procedure. It is well known that a good strategy for increasing n is to double it after every iteration. The criteria for thinning depend on the coupling used for constructing \mathbf{Z} . However, the dominating process \mathbf{D} depends only on λ . In summary, a dominated CFTP algorithm is described by Algorithm 5.1.

Consider the *backward coalescence* time $N^* = \min \{n \in \mathbb{N}_0 : \mathbf{L}_n(0) = \mathbf{U}_n(0)\}$. The average running time complexity of Algorithm 5.1 depends on the number of operations involved within N^* , which further depends on the construction of the interaction process and the bounding processes. At each iteration, the length of the dominating process \mathbf{D} is doubled on average backwards in time. Hence, on average the running time complexity doubles at each iteration. From the definition of N^* , the length of the last iteration is

$$2^{\lceil \log_2 N^* \rceil} \geq N^*.$$

Let $N^f = \min \{n \in \mathbb{N}_0 : \mathbf{L}_0(t_n) = \mathbf{U}_0(t_n)\}$ be the *forward* coalescence time. Because of the reversibility of the dominating process, it can be shown that N^* and N^f are identical in

Algorithm 5.1 Dominated CFTP

- 1: Generate $\{\mathbf{D}(t) : t_{-1} \leq t \leq 0\}$ with $\mathbf{D}(0) \sim \mu^0$
 - 2: $n \leftarrow 1$
 - 3: **repeat**
 - 4: $n \leftarrow 2 * n$
 - 5: Extend \mathbf{D} backwards from $\{\mathbf{D}(t) : t_{-n/2} \leq t \leq 0\}$ to $\{\mathbf{D}(t) : t_{-n} \leq t \leq 0\}$
 - 6: Construct $\{\mathbf{L}_n(t) : t_{-n} \leq t \leq 0\}$ and $\{\mathbf{U}_n(t) : t_{-n} \leq t \leq 0\}$
 - 7: **until** $\mathbf{L}_n(0) = \mathbf{U}_n(0)$
 - 8: **return** $\mathbf{L}_n(0)$
-

distribution [3], and hence the expected computational effort for constructing the dominating, upper bounding, and lower bounding processes up to the forward coalescence time N^f , starting from time 0, is a lower bound on the expected running time of the algorithm.

Below we consider three dominated CFTP methods applicable to the hard-sphere models.

5.1. Method 1

This method is based on [27]. Note that the *Papangelou conditional intensity* of the hard-sphere model is given by

$$\ell(\mathbf{x}, x) := \frac{I(\mathbf{x} \cup \{x\} \in \mathcal{A})}{I(\mathbf{x} \in \mathcal{A})} = I(\mathbf{x} \cup \{x\} \in \mathcal{A}), \tag{5.1}$$

with the convention that $0/0 = 0$. The interaction process $\mathbf{Z} = \{\mathbf{Z}(t) : t \in (-\infty, \infty)\}$ is constructed as follows. Suppose x is a birth to \mathbf{D} that sees \mathbf{Z} in a state $\mathbf{x} \in \mathcal{G}$. Then x is added to \mathbf{Z} if and only if $\ell(\mathbf{x}, x) = 1$. Every death in \mathbf{D} reflects in \mathbf{Z} ; that is, if there is a death of a point y in \mathbf{D} , then y is removed from the process \mathbf{Z} as well if it is present. It can be shown that μ is the unique invariant probability measure of \mathbf{Z} ; see, e.g., [16] or [12].

For each $n \geq 1$, the bounding processes are constructed as follows. As mentioned earlier, take $\mathbf{L}_n(t_{-n}) = \emptyset$ and $\mathbf{U}_n(t_{-n}) = \mathbf{D}(t_{-n})$. Suppose that $\mathbf{x}^l = \mathbf{L}_n(t_i)$ and $\mathbf{x}^u = \mathbf{U}_n(t_i)$ for $-n \leq i < 0$, then assign $\mathbf{L}_n(t) = \mathbf{x}^l$ and $\mathbf{U}_n(t) = \mathbf{x}^u$ for $t_i < t < t_{i+1}$. In case it is a birth x in the dominating process \mathbf{D} at time t_{i+1} , set $\mathbf{L}_n(t_{i+1}) = \mathbf{x}^l \cup \{x\}$ if $\ell(\mathbf{x}^u, x) = 1$; otherwise, it will remain unchanged, that is, $\mathbf{L}_n(t_{i+1}) = \mathbf{x}^l$. Similarly, set $\mathbf{U}_n(t_{i+1}) = \mathbf{x}^u \cup \{x\}$ if $\ell(\mathbf{x}^l, x) = 1$; otherwise, set $\mathbf{U}_n(t_{i+1}) = \mathbf{x}^u$. Every death in the dominating process reflects in both lower and upper bound processes. Note that a birth is accepted by the lower bounding process if the resulting state of the upper bounding process is in \mathcal{A} . Similarly, a birth in \mathbf{D} is accepted in the upper bounding process if the resulting state of the lower bounding process is in \mathcal{A} .

Theorem 5.1. *The expected running time complexity \mathcal{T}_{DC1} of the above dominated CFTP algorithm satisfies*

$$\mathcal{T}_{\text{DC1}} \geq c \frac{\lambda}{\mathcal{P}(\lambda)}, \tag{5.2}$$

for some constant $c > 0$. In particular,

$$\mathcal{T}_{DC1} = \begin{cases} \Omega(\lambda) & \text{if } \eta d \geq 2, \\ \Omega\left(\lambda \exp\left(\left(\frac{\gamma\mu_1}{2} + o(1)\right)\lambda^{2-\eta d}\right)\right) & \text{if } 1 < \eta d < 2, \\ \Omega\left(\lambda \exp\left(\left(1 + o(1)\right)\lambda\right)\right) & \text{if } 0 < \eta d \leq 1. \end{cases}$$

As highlighted by the numerical results in Section 6, the lower bound (5.2) is a loose bound, because it is established by considering the running time complexity only up to the time at which the lower bounding process receives its first arrival. This can be much smaller than the running time complexity up to the coalescence of the upper and lower bound processes.

5.2. Method 2

This method is an improved version of Method 1, again based on [27]. Observe that at any given time $t \in \mathbb{R}$, the interaction process $\mathbf{Z}(t)$ can have only non-overlapping spheres. This suggests a better way of constructing the bounding processes, which we now describe. For each $n \geq 1$, just as in Method 1, start with $\mathbf{L}_n(t_{-n}) = \emptyset$ and $\mathbf{U}_n(t_{-n}) = \mathbf{D}(t_{-n})$ to guarantee that $\mathbf{L}_n(t_{-n}) \subseteq \mathbf{Z}(t_{-n}) \subseteq \mathbf{U}_n(t_{-n})$. Suppose that the event at t_{i+1} is an arrival of sphere x . Irrespective of whether $\mathbf{U}_n(t_i) \in \mathcal{A}$ or not, if x does not overlap with any sphere in the upper bounding process $\mathbf{U}_n(t_i)$, then it cannot overlap with any sphere in $\mathbf{Z}(t_i)$, and hence it is accepted to \mathbf{Z} . Thus, we add x to both the bounding processes. (Observe that in Method 1, such an x is added to both the bounding processes only if $\mathbf{U}_n(t_i) \in \mathcal{A}$, because of the Papangelou conditional intensity (5.1).) If x overlaps with any sphere in the lower bounding process $\mathbf{L}_n(t_i)$, then it must overlap with a sphere in $\mathbf{Z}(t_i)$ as well, and hence it is not added to any of the bounding processes \mathbf{L} and \mathbf{U} . If x does not overlap with any sphere in $\mathbf{L}_n(t_i)$, but does overlap with a sphere in $\mathbf{U}_n(t_i)$, its presence in the process \mathbf{Z} cannot be ruled out, and hence we keep it in the upper bounding process, but not in the lower bounding process. Finally, every death in \mathbf{D} is reflected in both the bounding processes \mathbf{L}_n and \mathbf{U}_n . Under this construction, the lower bounding process accepts births more often, and hence the upper bounding process accepts births less often, compared with the construction in Section 5.1. As a result the running time of Method 2 is shorter than that of Method 1.

5.3. Method 3

A different approach for dominated CFTP for repulsive pairwise interaction processes has been proposed by Huber [22]. Here, we discuss the main ingredients of the method for hard-sphere models; refer to [22, 23] for more details. In this method, the interaction process \mathbf{Z} is different from the \mathbf{Z} in Sections 5.1–5.2; it is known as the spatial birth–death *swap* process, and its invariant distribution is again the distribution μ of the hard-sphere model. In addition to births and deaths of spheres, this process also allows *swap* moves; here a swap move is an event where an existing sphere is replaced by an arrival if it is the only sphere that overlaps with the arrival. The lower and upper bounding processes are constructed as follows. As usual, let $\mathbf{U}_n(t_{-n}) = \mathbf{D}(t_{-n})$ and $\mathbf{L}_n(t_{-n}) = \emptyset$. For any $0 < k < n$, if t_{-k} is the instant of a death in the dominating process $\mathbf{D}(t)$, then the death is reflected in both the upper and lower bound processes. Now suppose that $x \in \mathbf{D}(t_{-k})$ is born at t_{-k} .

Case 1: If no sphere in $U_n(t-k)$ overlaps with x , then the arrival sphere x is added to both $U_n(t-k)$ and $L_n(t-k)$. If only one sphere y in $U_n(t-k)$ overlaps with x , then y is removed from $U_n(t-k)$ (and from $L_n(t-k)$ if it is present), and x is added to both $U_n(t-k)$ and $L_n(t-k)$.

Case 2: There are at least two spheres in $L_n(t-k)$ overlapping with x . Then x is rejected by both $U_n(t-k)$ and $L_n(t-k)$.

Case 3: At most one sphere in $L_n(t-k)$ and at least two spheres in $U_n(t-k)$ overlap with x . Then x is added to $U_n(t-k)$ (but not to $L_n(t-k)$). If $y \in L_n(t-k)$ is the one that overlaps with x , then y is removed from $L_n(t-k)$.

6. Simulations

We compare the performance of all the methods discussed in this paper using numerical experiments, and illustrate the effectiveness of the proposed IS-based AR method in certain regimes where the other methods fail to work. For this, we consider the torus-hard-sphere model with a fixed radius \bar{r}/λ^η on the two-dimensional square $[0, 1]^2$. Thus, $\eta d = 2\eta$. In the first two experiments, by fixing values of η and \bar{r} , we estimate the complexity of each algorithm as a function of the intensity λ of the reference PPP by computing a sample average of the number of spheres (or, in this case, circles) generated per generation of a perfect sample of the hard-sphere model. Instead of estimating the expected running time complexities, we take this approach to keep the discussion independent of the underlying data structures and programming language used in the implementation of the algorithms. In addition, we estimate the non-overlapping probability $\mathcal{P}(\lambda)$ using the conditional Monte Carlo rare event estimation for *Gilbert graphs* proposed by [20]. The Gilbert graph under consideration is a random graph where the nodes constitute a λ -homogeneous PPP on $[0, 1]^2$ and there is an edge between two points if they are within a distance of $2\bar{r}/\lambda^\eta$. Therefore, $\mathcal{P}(\lambda)$ is the probability that there are no edges in the graph. The code for all the methods discussed in this paper is available at https://github.com/saratmoka/PerfectSampling_HardSpheres.

For the implementation of the proposed IS-based AR method, the grid is constructed using the cell-edge length $\epsilon = 1/\lfloor \lambda^\eta/\bar{r} \rfloor$; see Section 4.4 for more details on the selection of the cell-edge length. The complexities of the algorithms are estimated using 1000 samples. In the simulation results presented below, \widehat{S}_{NAR} and $\widehat{S}_{\text{ISAR}}$ denote the sample means of complexities of the naive AR and the IS-based AR algorithms. Likewise, $\widehat{S}_{\text{DCM1}}$, $\widehat{S}_{\text{DCM2}}$, and $\widehat{S}_{\text{DCM3}}$ respectively are the corresponding estimates for the three dominated CFTP methods (Methods 1, 2, and 3) presented in Section 5.

A standard program used for generating perfect samples of the hard-sphere model using dominated CFTP is `rHardcore()`, which is a part of the R package `Spatstat`, available at <https://spatstat.org>. Experiment 3 provides a perspective on the performance of the proposed method in relation to `rHardcore()` by comparing their expected running times as a function of \bar{r} for a fixed λ . We note that `rHardcore()` does not support the torus-hard-sphere model. However, when ‘`expand=TRUE`’ is selected, it reduces the boundary effects by generating a perfect sample on a larger window, and then clipping the result to the original window $[0, 1]^2$.

Experiment 1: In this experiment, we consider the high-density regime. Figure 3 compares the performance of all the algorithms for $\eta = 0.5$ (that is, $2\eta = 1$) and $\bar{r} = 1$ (this is identical to the regime where the underlying space is $[0, \sqrt{\lambda}]^2$ and the radius of each sphere is \bar{r}). This experiment suggests that the proposed IS-based AR method can perform significantly better than all the other methods. To show the rarity of the samples of the hard-sphere configurations under μ^0 , we plot $\log \mathcal{P}(\lambda)$ in Figure 4 and the expected intensity of the hard-sphere model in Figure 5.




FIGURE 3. Log of the expected number of points generated per perfect sample of the hard-sphere model, as a function of λ , in the regime where $\eta = 0.5$, $d = 2$, and $\bar{r} = 1$.




FIGURE 4. A plot of $\log \mathcal{P}(\lambda)$ versus λ in the regime where $\eta = 0.5$, $d = 2$, and $\bar{r} = 1$, where the non-overlapping probability $\mathcal{P}(\lambda)$ is estimated using the conditional Monte Carlo method proposed in [20]. The plot shows that, in the high-density regime, the configurations with hard spheres can be extremely rare under the measure μ^0 .

The significance of the proposed IS method in the high-density regime is more evident when $\eta = 0.25$ (that is, $2\eta = 0.5$) and $\bar{r} = 1$. In this case, for values of λ greater than 50, almost every time, the dominated CFTP algorithms terminated without producing any output. In particular, the `rHardcore` function terminated by producing the error ‘memory exhausted (limit reached?)’. On the other hand, the IS-based method generated 1000 samples in 0.13, 0.21, 68.94, 508.60, and 4315.19 seconds, respectively, for λ values of 50, 100, 200, 300, and 400.

Experiment 2: In this experiment, we consider the low-density regime. Figure 6 compares the performances of all the methods for $2\eta = 1.5$ and $\bar{r} = 0.5$ to illustrate the case where $1 < \eta d < 2$. As we can see, for large values of λ , the dominated CFTP methods 2 and 3 perform




FIGURE 5. The intensity of the hard-sphere model against λ in the regime where $\eta = 0.5$, $d = 2$, and $\bar{r} = 1$. Owing to the extreme rarity of the hard-sphere configurations under μ^0 , as shown in Figure 4, the intensity of the hard-sphere model is much smaller than the intensity λ of the PPP.




FIGURE 6. Log of the expected number of points generated per perfect sample of the hard-sphere model, as a function of λ , in the regime where $\eta = 0.75$, $d = 2$, and $\bar{r} = 0.5$.

better than the other methods, including the proposed method. Figure 7 is a plot of $\log \mathcal{P}(\lambda)$ against λ , while Figure 8 is a plot of the intensity of the hard-sphere model against λ .

Experiment 3: Figure 9 compares the running times of the proposed IS-based AR method and `rHardcore()` for generating 1000 samples. The same computer is used to run both programs. Here, we vary \bar{r} while fixing $\lambda = 50$ and $2\eta = 1$. Observe that for large values of \bar{r} the density is higher, and the proposed method performs far better than the dominated CFTP method. As we expect for this regime, as the radius \bar{r} increases, the intensity of the hard-sphere model decreases (Figure 11) while the rarity of the hard-sphere configurations under μ^0 increases (Figure 10).

7. Conclusion

In this paper we considered the problem of perfect sampling for Gibbs hard-sphere models on $[0, 1]^d$. We discussed the performance of the naive AR method and introduced IS-based




FIGURE 7. A plot of $\log \mathcal{P}(\lambda)$ versus λ in the regime where $\eta = 0.75$, $d = 2$, and $\bar{r} = 0.5$, where the non-overlapping probability $\mathcal{P}(\lambda)$ is estimated using the conditional Monte Carlo method proposed in [20]. Here we notice that the hard-sphere configurations are relatively less rare compared to the scenarios in Experiment 1.




FIGURE 8. The intensity of the hard-sphere model against λ in the regime where $\eta = 0.75$, $d = 2$, and $\bar{r} = 0.5$. Unlike in Experiment 1, the intensity of the hard-sphere model is relatively close to the intensity λ of the PPP.

enhancements to it. We also compared these methods to some of the popular CFTP-based techniques prevalent in the existing literature. For performance analysis and comparison (of expected running time complexity), we developed an asymptotic regime where the intensity λ of the reference PPP increased to infinity, while the (random) volume of each sphere, of the order of $\lambda^{-\eta d}$, decreased to zero, for different regimes of $\eta d > 0$. One main conclusion is that while the dominated CFTP methods perform better for $1 < \eta d < 2$ for large λ , our IS-based methods provide a significant improvement for $\eta d \leq 1$. En route, we established large deviations results for the probability that spheres do not overlap with each other when their centers constitute a PPP. We also conducted numerical experiments to validate our asymptotic results.

The proposed IS-based AR methods rely on cleverly partitioning the underlying configuration space and arriving at an appropriate change of measure on each partition. While we have shown how this may be effectively conducted for hard-sphere models, further research

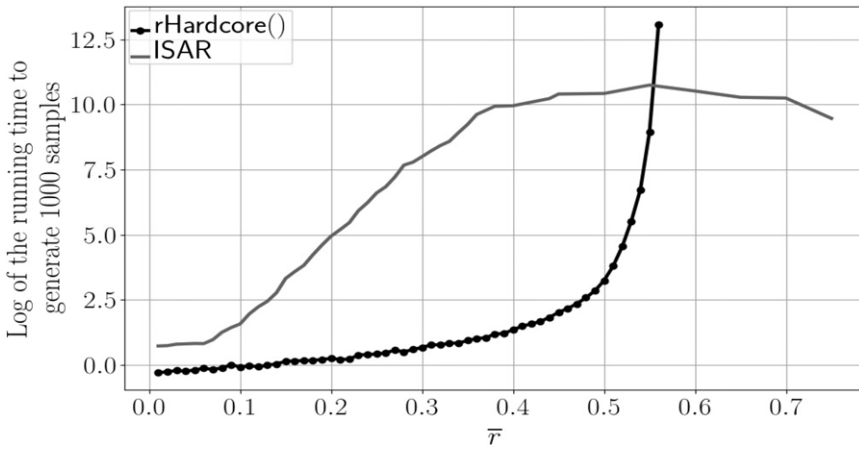


FIGURE 9. Comparison between the running times of the proposed IS-based AR method and `rHardcore()` for generating 1000 samples.




FIGURE 10. A plot of $\log \mathcal{P}(\lambda)$ versus \bar{r} in the regime where $\eta = 0.5$, $d = 2$, and $\lambda = 50$, where the non-overlapping probability $\mathcal{P}(\lambda)$ is again estimated using the conditional Monte Carlo method.




FIGURE 11. Intensity of the hard-sphere model against \bar{r} in the regime where $\eta = 0.5$, $d = 2$, and $\lambda = 50$.

is needed to develop effective implementations for perfect sampling from a broader class of Gibbs point processes.

Acknowledgements

We thank the referees and the associate editor for their comments and suggestions, which substantially improved the presentation of this paper. S. M. acknowledges the support of the Australian Research Council Centre of Excellence for Mathematical and Statistical Frontiers (ACEMS), under Grant No. CE140100049. S. J. acknowledges the support of the Department of Atomic Energy, Government of India, under Project No. RTI4001. M. M.’s research is partly funded by the NWO Gravitation project NETWORKS, Grant No. 024.002.003.

Appendix A. Proofs

The following lemmas are useful for proving Theorem 3, Proposition 4.1, and Proposition 4.3. Lemma A.1 is a standard Chernoff bound for Poisson random variables, and Lemma A.2 is Hoeffding’s inequality for U -statistics [21].

Lemma A.1. (Chernoff bound for Poisson) *Let $N \sim \text{Poi}(\lambda)$. Then, for any $0 < \epsilon < 1$,*

$$\mathbb{P}(N \leq (1 - \epsilon)\lambda) \leq \exp\left(-\frac{\lambda\epsilon^2}{2}\right) \quad \text{and} \quad \mathbb{P}(N \geq (1 + \epsilon)\lambda) \leq \exp\left(-\frac{\lambda\epsilon^2}{3}\right).$$

Lemma A.2. (Hoeffding, 1963) *Suppose that $\xi_1, \xi_2, \dots, \xi_n$ are i.i.d. random variables and $g: \mathbb{R}^k \rightarrow [0, 1]$ is a measurable function. Set*

$$Z_n = \sum_{1 \leq i_1 < i_2 < \dots < i_k \leq n} g(\xi_{i_1}, \xi_{i_2}, \dots, \xi_{i_k})$$

for a positive integer $k \leq n$ (this is known as a U -statistic of order k). Then, for any $\epsilon > 0$,

$$\mathbb{P}\left(Z_n \geq \frac{n}{k} \left(\mathbb{E}[g(\xi_1, \xi_2, \dots, \xi_k)] + \epsilon\right)\right) \leq 2 \exp\left(-2\lfloor n/k \rfloor \epsilon^2\right).$$

The same estimate holds for $\mathbb{P}\left(Z_n \leq \frac{n}{k} \left(\mathbb{E}[g(\xi_1, \xi_2, \dots, \xi_k)] - \epsilon\right)\right)$.

A.1. Proof of Theorem 3.1

Recall that $\lambda > 0, \eta > 0$, and $\frac{R_1}{\lambda^\eta}, \dots, \frac{R_n}{\lambda^\eta}$ are the radii of n spheres whose respective centers Y_1, \dots, Y_n are independently and uniformly generated on the d -dimensional unit cube $[0, 1]^d$, where R_1, \dots, R_n are i.i.d. positive random variables bounded from above by a constant \bar{r} and are independent of Y_1, \dots, Y_n . Define $m_i := \mathbb{E}[(R_1 + R_2)^{id}]$, for all $i = 1, 2, \dots$. Let $\mathcal{P}_n(\lambda)$ be the probability that these n spheres do not overlap with each other. Since the number of spheres in a λ -homogeneous marked PPP on $[0, 1]^d$ is a Poisson random variable with mean λ , the non-overlapping probability is

$$\mathcal{P}(\lambda) = \mathbb{E}[\mathcal{P}_N(\lambda)], \tag{A.1}$$

where $N \sim \text{Poi}(\lambda)$.

In proving the theorem, we use the following lemmas, which exploit the reference IS measure $\tilde{\mu}_n$ introduced in Section 4.2. By (4.7) and the definition of $\mathcal{P}_n(\lambda)$,

$$\begin{aligned} \mathcal{P}_n(\lambda) &= \mathbb{P}_{\mu^0}(\mathbf{X} \in \mathcal{A} | |\mathbf{X}| = n) \\ &= \mathbb{E}_{\tilde{\mu}_n} \left[I(\mathbf{X} \in \mathcal{A}) \prod_{i=1}^n (1 - B_i) \right] \\ &= \mathbb{E}_{\tilde{\mu}_n} \left[\prod_{i=1}^n (1 - B_i) \right]. \end{aligned} \tag{A.2}$$

The following bound holds trivially:

$$B_i \leq \frac{\gamma}{\lambda^{\eta d}} \sum_{j=1}^{i-1} (R_j + R_i)^d, \tag{A.3}$$

where the sum is taken to be zero when $i = 1$. Let

$$\theta_{n,\lambda} = \frac{\gamma(2\bar{r})^d n}{\lambda^{\eta d}}.$$

We have the following upper and lower bounds on $\mathcal{P}_n(\lambda)$.

Lemma A.3. *For the above set-up,*

$$\mathcal{P}_n(\lambda) \geq \exp \left(-n \sum_{j=1}^{\infty} \left(\frac{\gamma n}{\lambda^{\eta d}} \right)^j \frac{m_j}{j(j+1)} \right). \tag{A.4}$$

Furthermore, for any $\epsilon > 0$,

$$\mathcal{P}_n(\lambda) \leq N_{n,\lambda} \left[\exp \left(-\frac{\gamma n(n-1)(m_1 - \epsilon)}{2\lambda^{\eta d}} \right) + 2 \exp \left(-\frac{(n-1)\epsilon^2}{(2\bar{r})^{2d}} \right) \right] \tag{A.5}$$

for any n and λ such that $\theta_{n,\lambda} < 1$, where $N_{n,\lambda}$ is a function of n, λ , and \bar{r} , defined by (A.8) in the proof below, such that

$$\lim_{\lambda \rightarrow \infty} \frac{1}{\lambda^{2-\eta d}} \log N_{\lambda,\lambda} = 0 \quad \text{if } \eta d > 1. \tag{A.6}$$

In particular, for the torus-hard-sphere model,

$$\lim_{\lambda \rightarrow \infty} N_{\lambda,\lambda} = 1 \quad \text{if } \eta d > 3/2. \tag{A.7}$$

Proof. Lower bound: To prove (A.4), notice that, by (A.2),

$$\mathcal{P}_n(\lambda) = \mathbb{E}_{\tilde{\mu}_n} \left[\exp \left(\sum_{i=1}^n \log (1 - B_i) \right) \right] = \mathbb{E}_{\tilde{\mu}_n} \left[\exp \left(\sum_{i=1}^n \sum_{j=1}^{\infty} -\frac{1}{j} B_i^j \right) \right],$$

using the Taylor expansion $\log(1 - x) = -\sum_{j=1}^{\infty} x^j/j$ for $0 \leq x \leq 1$. By Jensen's inequality and (A.3),

$$\begin{aligned} \mathcal{P}_n(\lambda) &\geq \exp\left(-\sum_{i=1}^n \sum_{j=1}^{\infty} \frac{1}{j} \mathbb{E}_{\tilde{\mu}_n} [B_i^j]\right) \\ &\geq \exp\left(-\sum_{i=1}^n \sum_{j=1}^{\infty} \frac{\gamma^j}{j \lambda^{j \eta d}} \mathbb{E}\left[\left(\sum_{l=1}^{i-1} (R_l + R_i)^d\right)^j\right]\right) \\ &\geq \exp\left(-\sum_{i=1}^n \sum_{j=1}^{\infty} \frac{\gamma^j (i-1)^j}{j \lambda^{j \eta d}} \mathbb{E}\left[\left(\frac{1}{i-1} \sum_{l=1}^{i-1} (R_l + R_i)^d\right)^j\right]\right). \end{aligned}$$

Again by Jensen's inequality,

$$\left(\frac{1}{i-1} \sum_{l=1}^{i-1} (R_l + R_i)^d\right)^j \leq \frac{1}{i-1} \left(\sum_{l=1}^{i-1} (R_l + R_i)^{jd}\right),$$

and thus

$$\mathcal{P}_n(\lambda) \geq \exp\left(-\sum_{j=1}^{\infty} \frac{\gamma^j m_j}{j \lambda^{j \eta d}} \sum_{i=1}^n (i-1)^j\right).$$

We establish (A.4) using

$$\sum_{i=1}^n (i-1)^j \leq \int_{x=0}^n x^j dx = \frac{n^{j+1}}{j+1}.$$

Upper bound: Let $R_{(1)}, R_{(2)}, \dots, R_{(n)}$ be the order statistics of R_1, R_2, \dots, R_n . Since the non-overlapping probability $\mathcal{P}_n(\lambda)$ is independent of the order in which the spheres are generated, without loss of generality assume that the i th sphere has radius $R_{(i)}$. For each $1 \leq j \leq i-1$, let $\tilde{B}_i(j)$ be the volume of the blocked region for the i th sphere generation when the $(j+1)$ th, $(j+2)$ th, \dots , $(i-1)$ th spheres are ignored, where $\tilde{B}_i(0) = 0$. We can think of $\tilde{B}_i(j) - \tilde{B}_i(j-1)$ as the blocking volume contributed by the j th sphere for the i th sphere. Under the new measure $\tilde{\mu}_n$, the blocking volume seen by the i th sphere is $B_i = \sum_{j=1}^{i-1} (\tilde{B}_i(j) - \tilde{B}_i(j-1))$. Consider the sets

$$\mathcal{A}^{(i)} := \left\{j \in \{1, 2, \dots, i-1\} : \tilde{B}_i(j) - \tilde{B}_i(j-1) = \frac{\gamma}{\lambda \eta d} (R_{(j)} + R_{(i)})^d\right\},$$

for $i \leq n$, and take $\bar{\mathcal{N}}^{(i)} := \{1, 2, \dots, i-1\} \setminus \mathcal{A}^{(i)}$. Using the inequality $1 - x \leq e^{-x}$ and (A.2),

$$\begin{aligned} \mathcal{P}_n(\lambda) &\leq \mathbb{E}_{\tilde{\mu}_n} \left[\exp\left(-\sum_{i=1}^n B_i\right) \right] = \mathbb{E}_{\tilde{\mu}_n} \left[\exp\left(-\sum_{i=1}^n \sum_{j=1}^{i-1} (\tilde{B}_i(j) - \tilde{B}_i(j-1))\right) \right] \\ &\leq \mathbb{E}_{\tilde{\mu}_n} \left[\exp\left(-\frac{\gamma}{\lambda \eta d} \sum_{i=1}^n \sum_{j \in \bar{\mathcal{N}}^{(i)}} (R_{(j)} + R_{(i)})^d\right) \right] \end{aligned}$$

$$\begin{aligned}
 &= \mathbb{E}_{\tilde{\mu}_n} \left[\exp \left(-\frac{\gamma (2\bar{r})^d}{\lambda^{\eta d}} Z_n + \frac{\gamma}{\lambda^{\eta d}} \sum_{i=1}^n \sum_{j \in \bar{\mathcal{N}}^{(i)}} (R_{(j)} + R_{(i)})^d \right) \right] \\
 &\leq \mathbb{E}_{\tilde{\mu}_n} \left[\exp \left(-\frac{\gamma (2\bar{r})^d}{\lambda^{\eta d}} Z_n + \frac{\gamma (2\bar{r})^d}{\lambda^{\eta d}} \sum_{i=1}^n |\bar{\mathcal{N}}^{(i)}| \right) \right],
 \end{aligned}$$

where

$$Z_n = \sum_{i=1}^n \sum_{j=1}^{i-1} \left(\frac{R_{(j)} + R_{(i)}}{2\bar{r}} \right)^d,$$

and the last inequality holds from the assumption that each $R_i \leq \bar{r}$. Since $R_{(i)}$ is non-decreasing with i , from the definition of $\mathcal{N}^{(i)}$, it is easy to see that $|\bar{\mathcal{N}}^{(i)}|$ is non-decreasing with i . Therefore,

$$\begin{aligned}
 \mathcal{P}_n(\lambda) &\leq \mathbb{E}_{\tilde{\mu}_n} \left[\exp \left(-\frac{\gamma (2\bar{r})^d}{\lambda^{\eta d}} Z_n + \theta_{n,\lambda} |\bar{\mathcal{N}}^{(n)}| \right) \right] \\
 &= \mathbb{E}_{\tilde{\mu}_n} \left[\exp \left(-\frac{\gamma (2\bar{r})^d}{\lambda^{\eta d}} Z_n \right) \mathbb{E}_{\tilde{\mu}_n} \left[\exp \left(\theta_{n,\lambda} |\bar{\mathcal{N}}^{(n)}| \right) \mid R_1, \dots, R_n \right] \right].
 \end{aligned}$$

We now show that $|\bar{\mathcal{N}}^{(n)}|$ is stochastically bounded by a binomial random variable, and as a consequence, the conditional expectation

$$\mathbb{E}_{\tilde{\mu}_n} \left[\exp \left(\theta_{n,\lambda} |\bar{\mathcal{N}}^{(n)}| \right) \mid R_1, \dots, R_n \right]$$

is uniformly bounded from above by a constant, which is a function of n , λ , and \bar{r} . Let

$$q_j = \mathbb{P}_{\tilde{\mu}_n} \left(\tilde{B}_n(j) - \tilde{B}_n(j-1) < \frac{\gamma}{\lambda^{\eta d}} (R_{(j)} + R_{(n)})^d \right).$$

Clearly, q_j is increasing with j , and therefore $q_j \leq q_{n-1}$ for all $j \leq n-1$. This implies that $|\bar{\mathcal{N}}^{(n)}|$ is stochastically bounded by a binomial random variable with parameters n and q_{n-1} , and thus

$$\mathbb{E}_{\tilde{\mu}_n} \left[\exp \left(\theta_{n,\lambda} |\bar{\mathcal{N}}^{(n)}| \right) \mid R_1, \dots, R_n \right] \leq \left(q_{n-1} \exp(\theta_{n,\lambda}) + (1 - q_{n-1}) \right)^n.$$

Because of the boundary effect, q_{n-1} is not the same for the torus model (where boundary spheres loop over to the opposite boundaries) and the Euclidean model. Observe that, for the Euclidean model,

$$\tilde{B}_n(n-1) - \tilde{B}_n(n-2) < \frac{\gamma (R_{(j)} + R_{(n)})^d}{\lambda^{\eta d}}$$

if either

- (1) the center of the $(n-1)$ th sphere is within a distance of $(R_{(j)} + R_{(n-1)} + 2R_{(n)})/\lambda^\eta$ from the center of the j th sphere for some $j \leq n-2$, or
- (2) the center of $(n-1)$ th sphere is within a distance of $(R_{(n-1)} + R_{(n)})/\lambda^\eta$ from the boundary of the unit cube.

Note that the boundary event (2) is irrelevant for the torus-hard-sphere model. The probability of the event (1) is maximized by

$$\frac{\gamma}{\lambda^{\eta d}} \sum_{j=1}^{n-2} \frac{(R_{(j)} + R_{(n-1)} + 2R_{(n)})^d}{1 - B_{n-1}},$$

while that of the event (2) is maximized by

$$\frac{1 - (1 - 2(R_{(n-1)} + R_{(n)})/\lambda^\eta)^d}{1 - B_{n-1}}.$$

Since the R_i are bounded from above by \bar{r} and $B_{n-1} \leq \theta_{n,\lambda}$ (from (A.3)), we have

$$q_{n-1} \leq \bar{q}_{n,\lambda} := \begin{cases} \frac{2^d \theta_{n,\lambda}}{1 - \theta_{n,\lambda}} + \frac{c}{\lambda^\eta (1 - \theta_{n,\lambda})} & \text{for the Euclidean model,} \\ \frac{2^d \theta_{n,\lambda}}{1 - \theta_{n,\lambda}} & \text{for the torus model,} \end{cases}$$

for any n and λ such that $\theta_{n,\lambda} < 1$, and for some constant c . Let

$$N_{n,\lambda} = \left(1 + \bar{q}_{n,\lambda} \left(\exp(\theta_{n,\lambda}) - 1\right)\right)^n; \tag{A.8}$$

then

$$\mathcal{P}_n(\lambda) \leq N_{n,\lambda} \mathbb{E}_{\tilde{\mu}_n} \left[\exp\left(-\frac{\gamma (2\bar{r})^d}{\lambda^{\eta d}} Z_n\right) \right].$$

Using the definition of Z_n , for any $\epsilon > 0$,

$$\mathbb{E}_{\tilde{\mu}_n} \left[\exp\left(-\frac{\gamma (2\bar{r})^d}{\lambda^{\eta d}} Z_n\right) \right] \leq \exp\left(-\frac{\gamma n(n-1)(m_1 - \epsilon)}{2\lambda^{\eta d}}\right) + \mathbb{P}_{\tilde{\mu}_n} \left(Z_n < \frac{n(n-1)(m_1 - \epsilon)}{2(2\bar{r})^d} \right).$$

By Lemma 8 (with $k = 2$),

$$\mathbb{P}_{\tilde{\mu}_n} \left(Z_n < \frac{n(n-1)}{2} \left(\frac{m_1 - \epsilon}{(2\bar{r})^d} \right) \right) \leq 2 \exp\left(-\frac{(n-1)\epsilon^2}{(2\bar{r})^{2d}}\right),$$

and thus (A.5) is established.

It remains to prove (A.6) and (A.7) under the assumption that $\eta d > 1$. For this case, we have that $\lim_{\lambda \nearrow \infty} \theta_{\lambda,\lambda} = 0$ and hence $\lim_{\lambda \nearrow \infty} \bar{q}_{\lambda,\lambda} = 0$. Since

$$N_{\lambda,\lambda} \leq \exp\left(\lambda \bar{q}_{\lambda,\lambda} \left[\exp(\theta_{\lambda,\lambda}) - 1\right]\right),$$

using the Taylor expansion of the exponential function,

$$0 \leq \lim_{\lambda \nearrow \infty} \frac{1}{\lambda^{2-\eta d}} \log N_{\lambda,\lambda} \leq \lim_{\lambda \nearrow \infty} \left[\bar{q}_{\lambda,\lambda} \sum_{j \in \mathbb{N}_0} \frac{\gamma^{(j+1)} (2\bar{r})^{(j+1)d}}{(j+1)!} \lambda^{j(\eta d - 1)} \right] = 0.$$

Thus (A.6) holds. In particular, for the torus model with $\eta d > 3/2$,

$$\lambda \bar{q}_{\lambda,\lambda} \left[\exp(\theta_{\lambda,\lambda}) - 1\right] = \lambda \frac{2^d \theta_{\lambda,\lambda}}{1 - \theta_{\lambda,\lambda}} \left[\exp(\theta_{\lambda,\lambda}) - 1\right] = \frac{2^d}{1 - \theta_{\lambda,\lambda}} \sum_{j=2}^{\infty} \frac{\gamma^j (2\bar{r})^{jd}}{j!} \lambda^{1-j(\eta d - 1)}$$

goes to 0 as $\lambda \nearrow \infty$, and hence (A.7) holds. □

Lemma A.4. Suppose that $1 < \eta d \leq 2$. Then, for any $0 < a < 0.5$,

$$\mathcal{P}(\lambda) \geq \exp \left(- \sum_{j=1}^{\infty} \frac{\lambda^{j(1-\eta d)+1} \left(1 + \frac{1}{\lambda^a}\right)^{j+1} \gamma^j m_j}{j(j+1)} \right) [1 - o(1)]. \quad (\text{A.9})$$

Furthermore, let $\bar{\lambda} = \lceil \lambda(1 - \frac{1}{\lambda^a}) \rceil$ for some constant a such that $0 < a < \frac{\eta d - 1}{2}$. Then, for any $\epsilon > 0$,

$$\mathcal{P}(\lambda) \leq N_{\lambda, \lambda} \exp \left(- \frac{\gamma \bar{\lambda}^2 (m_1 - \epsilon)}{2\lambda^{\eta d}} \right) [1 + o(1)], \quad (\text{A.10})$$

where $N_{\lambda, \lambda}$ satisfies (A.6) and (A.7). In particular, (A.9) holds with $\epsilon = 0$ if $\eta d > 5/3$ and $2 - \eta d < a < \frac{\eta d - 1}{2}$.

Proof. Lower bound: Fix a such that $0 < a < 0.5$. Since $\mathcal{P}_n(\lambda)$ is a decreasing function of n for any fixed λ , by Lemma 8.1, we can say that for all $n < \lambda \left(1 + \frac{1}{\lambda^a}\right)$,

$$\mathcal{P}_n(\lambda) \geq \exp \left(- \sum_{j=1}^{\infty} \frac{\lambda^{j(1-\eta d)+1} \left(1 + \frac{1}{\lambda^a}\right)^{j+1} \gamma^j m_j}{j(j+1)} \right),$$

and from (A.1) and the Chernoff bound for the Poisson variable N (see Lemma A.1),

$$\begin{aligned} \mathcal{P}(\lambda) &\geq \mathbb{E} \left(\mathcal{P}_N(\lambda); N < \lambda \left(1 + \frac{1}{\lambda^a}\right) \right) \\ &\geq \mathbb{P} \left(N < \lambda \left(1 + \frac{1}{\lambda^a}\right) \right) \exp \left(- \sum_{j=1}^{\infty} \frac{\lambda^{j(1-\eta d)+1} \left(1 + \frac{1}{\lambda^a}\right)^{j+1} \gamma^j m_j}{j(j+1)} \right) \\ &\geq \left(1 - \exp \left(-\frac{1}{3} \lambda^{1-2a} \right) \right) \exp \left(- \sum_{j=1}^{\infty} \frac{\lambda^{j(1-\eta d)+1} \left(1 + \frac{1}{\lambda^a}\right)^{j+1} \gamma^j m_j}{j(j+1)} \right). \end{aligned}$$

Now (A.9) follows easily because $\exp \left(-\frac{1}{3} \lambda^{1-2a} \right) = o(1)$ as a function of λ .

Upper bound: From (A.1),

$$\mathcal{P}(\lambda) = \mathbb{E}[\mathcal{P}_N(\lambda)] \leq \mathbb{E}[\mathcal{P}_N(\lambda); N \geq \bar{\lambda}] + \mathbb{P}(N < \bar{\lambda}) \leq \mathcal{P}_{\bar{\lambda}}(\lambda) + \mathbb{P}(N < \bar{\lambda}), \quad (\text{A.11})$$

where the last inequality is due to the fact that $\mathcal{P}_n(\lambda)$ is a decreasing function of n for fixed λ . We now analyze $\mathcal{P}_{\bar{\lambda}}(\lambda)$ and $\mathbb{P}(N < \bar{\lambda})$ separately.

By Lemma A.3, for any $\epsilon > 0$,

$$\begin{aligned} \mathcal{P}_{\bar{\lambda}}(\lambda) &\leq N_{\lambda, \lambda} \left[\exp \left(- \frac{\gamma \bar{\lambda} (\bar{\lambda} - 1) (m_1 - \epsilon)}{2\lambda^{\eta d}} \right) + 2 \exp \left(- \frac{\bar{\lambda} \epsilon^2}{(2\bar{r})^{2d}} \right) \right] \\ &\leq N_{\lambda, \lambda} \left[\exp \left(- \frac{\gamma \bar{\lambda}^2 (m_1 - \epsilon)}{2\lambda^{\eta d}} \right) \exp \left(\frac{\gamma m_1}{2\lambda^{\eta d - 1}} \right) + 2 \exp \left(- \frac{\bar{\lambda} \epsilon^2}{(2\bar{r})^{2d}} \right) \right], \end{aligned}$$

where we used the fact that $\bar{\lambda} \leq \lambda$. We rewrite the above expression as follows:

$$\mathcal{P}_{\bar{\lambda}}(\lambda) \leq N_{\lambda,\lambda} \exp\left(-\frac{\gamma \bar{\lambda}^2 (m_1 - \epsilon)}{2\lambda^{\eta d}}\right) \left(\exp\left(\frac{\gamma m_1}{2\lambda^{\eta d-1}}\right) + 2 \exp\left(\frac{\gamma \bar{\lambda}^2 m_1}{2\lambda^{\eta d}} - \frac{\bar{\lambda} \epsilon^2}{(2\bar{r})^{2d}}\right) \right).$$

Note that

$$\frac{\gamma \bar{\lambda}^2 m_1}{2\lambda^{\eta d}} = O\left(\lambda^{2-\eta d}\right)$$

and

$$\frac{\bar{\lambda} \epsilon^2}{(2\bar{r})^{2d}} = \Omega(\lambda).$$

Since $\eta d > 1$,

$$2 \exp\left(\frac{\gamma \bar{\lambda}^2 m_1}{2\lambda^{\eta d}} - \frac{\bar{\lambda} \epsilon^2}{(2\bar{r})^{2d}}\right) \leq 2 \exp\left(-\bar{\lambda} \left(\frac{\epsilon^2}{(2\bar{r})^{2d}} - \frac{\gamma m_1}{2\lambda^{\eta d-1}}\right)\right) \rightarrow 0 \quad \text{as } \lambda \rightarrow \infty, \quad (\text{A.12})$$

and since

$$\lim_{\lambda \rightarrow \infty} \exp\left(\frac{\gamma m_1}{2\lambda^{\eta d-1}}\right) = 1,$$

we can say that the first term $\mathcal{P}_{\bar{\lambda}}(\lambda)$ in (A.11) satisfies the following inequality:

$$\mathcal{P}_{\bar{\lambda}}(\lambda) \leq N_{\lambda,\lambda} \exp\left(-\frac{\gamma \bar{\lambda}^2 (m_1 - \epsilon)}{2\lambda^{\eta d}}\right) [1 + o(1)].$$

By Lemma A.1,

$$\mathbb{P}(N \leq \bar{\lambda}) \leq \mathbb{P}\left(N \leq \lambda \left(1 - \frac{1}{\lambda^a}\right)\right) \leq \exp\left(-\frac{\lambda^{1-2a}}{2}\right). \quad (\text{A.13})$$

Since $2a < 1$ (because $\eta d \leq 2$), we have $1 - 2a > 2 - \eta d$, and hence, using (A.13) and the fact that $N_{\lambda,\lambda} \geq 1$,

$$\frac{\exp\left(\frac{\gamma \bar{\lambda}^2 (m_1 - \epsilon)}{\lambda^{\eta d}}\right) \mathbb{P}(N \leq \bar{\lambda})}{N_{\lambda,\lambda}} \leq \exp\left(\frac{\gamma \bar{\lambda}^2 (m_1 - \epsilon)}{2\lambda^{\eta d}} - \frac{\lambda^{1-2a}}{2}\right) \rightarrow 0 \quad \text{as } \lambda \rightarrow \infty. \quad (\text{A.14})$$

Therefore (A.10) follows from (A.11) and (A.14).

In particular if $\eta d > 5/3$, we choose a such that $2 - \eta d < a < \frac{\eta d - 1}{2}$. Let $\epsilon = 1/\lambda^a$. Then (A.12) and (A.14) hold. We complete the proof using the fact that

$$\lim_{\lambda \rightarrow \infty} \exp\left(\frac{\gamma \bar{\lambda}^2 \epsilon}{2\lambda^{\eta d}}\right) = 1. \quad \square$$

Proof of Theorem 3.1. The following upper and lower bounds together complete the proof.

Lower bounds: Consider the inequality (A.9).

Case: $\eta d > 2$. Since $R \leq \bar{r}$, we have $m_j \leq (2\bar{r})^{jd}$. Thus, for $\eta d > 2$, all the terms in the exponent of the right-hand side of (A.9) go to zero asymptotically. In other words, for any $0 < a < 0.5$,

$$\lim_{\lambda \rightarrow \infty} \sum_{j=1}^{\infty} \frac{\lambda^{j(1-\eta d)+1} \left(1 + \frac{1}{\lambda^a}\right)^{j+1} \gamma^j m_j}{j(j+1)} = 0.$$

This means that $\lim_{\lambda \rightarrow \infty} \mathcal{P}(\lambda) = 1$ for $\eta d > 2$.

Case: $3/2 < \eta d \leq 2$. Using (A.9),

$$\mathcal{P}(\lambda) \exp\left(\frac{\gamma m_1}{2} \lambda^{2-\eta d}\right)$$

is bounded from below by

$$\exp\left(O\left(\lambda^{2-\eta d-a}\right) - \sum_{j=2}^{\infty} \frac{\lambda^{j(1-\eta d)+1} \left(1 + \frac{1}{\lambda^a}\right)^{j+1} \gamma^j m_j}{j(j+1)}\right) [1 - o(1)].$$

By fixing $a > 2 - \eta d$, we see that the right-hand side of the above expression goes to 1 as $\lambda \nearrow \infty$. Thus,

$$\liminf_{\lambda \rightarrow \infty} \left[\mathcal{P}(\lambda) \exp\left(\frac{\gamma m_1}{2} \lambda^{2-\eta d}\right)\right] \geq 1.$$

Case: $1 < \eta d \leq 3/2$. By applying log on both sides of (A.9), we have for any $0 < a < 0.5$ that

$$\log \mathcal{P}(\lambda) \geq \log(1 - o(1)) - \sum_{j=1}^{\infty} \frac{\lambda^{j(1-\eta d)+1} \left(1 + \frac{1}{\lambda^a}\right)^{j+1} \gamma^j m_j}{j(j+1)},$$

and we see that

$$\frac{1}{\lambda^{2-\eta d}} \sum_{j=1}^{\infty} \frac{\lambda^{j(1-\eta d)+1} \left(1 + \frac{1}{\lambda^a}\right)^{j+1} \gamma^j m_j}{j(j+1)} = \frac{\left(1 + \frac{1}{\lambda^a}\right)^2 \gamma m_1}{2} + \sum_{j=2}^{\infty} \frac{\left(1 + \frac{1}{\lambda^a}\right)^{j+1} \gamma^j m_j}{\lambda^{(j-1)(\eta d-1)} j(j+1)}.$$

Thus

$$\liminf_{\lambda \rightarrow \infty} \frac{1}{\lambda^{2-\eta d}} \log \mathcal{P}(\lambda) \geq -\frac{\gamma m_1}{2}$$

for $\eta d > 1$.

Case: $0 < \eta d \leq 1$. Configurations with one sphere or no sphere are always accepted; that is, $\mathcal{P}_1(\lambda) = \mathcal{P}_0(\lambda) = 1$. The probability of generating no sphere is $e^{-\lambda}$. Consequently, $\mathcal{P}(\lambda) > e^{-\lambda}$, and for any $\eta d > 0$,

$$\liminf_{\lambda \rightarrow \infty} \left[\frac{1}{\lambda} \log \mathcal{P}(\lambda)\right] \geq -1. \tag{A.15}$$

In particular, assume that $\eta d = 1$. For this case, first we show that the limit $\delta := \lim_{\lambda \rightarrow \infty} \left[\frac{1}{\lambda} \log \mathcal{P}(\lambda)\right]$ exists. To prove this, partition the cube $[0, 1]^d$ into a cubic grid of cell-edge length $x^{1/d} \in (0, 1)$. Ignore the cells at the boundary that have edge length strictly smaller than $x^{1/d}$. The total intensity of the underlying PPP over a cell is then λx .

When $\eta d = 1$, the radius of each sphere is identical in distribution to R/λ . Observe that the non-overlapping probability of the spheres restricting to a cell is $\mathcal{P}(\lambda x)$ (see the definition of the non-overlapping probability). Since the total number of cells is at least $1/x$, the non-overlapping probability $\mathcal{P}(\lambda) \leq (\mathcal{P}(\lambda x))^{1/x}$, and thus $\frac{1}{\lambda} \log \mathcal{P}(\lambda) \leq \frac{1}{\lambda x} \log \mathcal{P}(\lambda x)$. We can

increase λ and decrease the cell-edge length $x^{1/d}$ so that $y := \lambda x$ is fixed. Then the following inequality holds:

$$\limsup_{\lambda \rightarrow \infty} \left[\frac{1}{\lambda} \log \mathcal{P}(\lambda) \right] \leq \frac{1}{y} \log \mathcal{P}(y) < 0.$$

Now the existence of the required limit is established by applying the limit on y :

$$\limsup_{\lambda \rightarrow \infty} \left[\frac{1}{\lambda} \log \mathcal{P}(\lambda) \right] \leq \liminf_{y \rightarrow \infty} \left[\frac{1}{y} \log \mathcal{P}(y) \right].$$

To show that $\delta \nearrow 0$ as $\gamma m_1 \searrow 0$, assume that $\gamma m_1 < \epsilon$ for a constant $\epsilon \in (0, 1)$. By (A.2) and (A.3),

$$\mathcal{P}_n(\lambda) \geq \mathbb{E} \left[\prod_{i=1}^{n-1} \left(1 - \frac{\gamma}{\lambda} \sum_{k=1}^i (R_k + R_i)^d \right)^+ \right].$$

Consider the following partial order on \mathbb{R}_+^n : for any $y, y' \in \mathbb{R}_+^n$, we say that $y \leq y'$ if $y_i \leq y'_i$ for all $i = 1, \dots, n$. A function $f: \mathbb{R}_+^n \rightarrow \mathbb{R}$ is called *increasing* (or *decreasing*) if $f(y) \leq f(y')$ (or $f(y) \geq f(y')$) for all $y, y' \in \mathbb{R}_+^n$ such that $y \leq y'$. If f and g are either increasing or decreasing functions, then Theorem 2.4 of [17] (the Fortuin–Kasteleyn–Ginibre (FKG) inequality) can be trivially extended to show that $\mathbb{E}[f(Y)g(Y)] \geq \mathbb{E}[f(Y)]\mathbb{E}[g(Y)]$. Clearly the following function f_i is a decreasing function on \mathbb{R}_+^n :

$$f_i(y) = \left(1 - \frac{\gamma}{\lambda} \sum_{k=1}^i (y_k + y_i)^d \right)^+.$$

Therefore,

$$\mathbb{E} \left[\prod_{i=1}^{n-1} \left(1 - \frac{\gamma}{\lambda} \sum_{k=1}^i (R_k + R_i)^d \right)^+ \right] \geq \prod_{i=1}^{n-1} \mathbb{E} \left[\left(1 - \frac{\gamma}{\lambda} \sum_{k=1}^i (R_k + R_i)^d \right)^+ \right].$$

Using the convexity of the function x^+ and Jensen’s inequality, for each i ,

$$\mathbb{E} \left[\left(1 - \frac{\gamma}{\lambda} \sum_{k=1}^i (R_k + R_i)^d \right)^+ \right] \geq \left(1 - i \frac{\gamma m_1}{\lambda} \right)^+,$$

and thus

$$\mathcal{P}_n(\lambda) \geq \prod_{i=1}^{n-1} \left(1 - i \frac{\gamma m_1}{\lambda} \right)^+.$$

With $\underline{\lambda} = \lfloor \lambda + \lambda^{0.75} \rfloor$ and $N \sim Poi(\lambda)$,

$$\mathcal{P}(\lambda) = \sum_{n=0}^{\infty} e^{-\lambda} \frac{\lambda^n}{n!} \mathcal{P}_n(\lambda) \geq \sum_{n=0}^{\underline{\lambda}} e^{-\lambda} \frac{\lambda^n}{n!} \mathcal{P}_n(\lambda) \geq \mathcal{P}_{\underline{\lambda}}(\lambda) \mathbb{P}(N \leq \underline{\lambda}).$$

By applying log on both the sides of the above inequality and scaling with $1/\lambda$, we obtain

$$\frac{1}{\lambda} \log \mathcal{P}(\lambda) \geq \frac{1}{\lambda} \log \mathcal{P}_{\underline{\lambda}}(\lambda) + \frac{1}{\lambda} \log \mathbb{P} (N \leq \underline{\lambda}) .$$

From the definition of $\underline{\lambda}$ and Lemma A.1, the second term, $\frac{1}{\lambda} \log \mathbb{P} (N \leq \underline{\lambda})$, goes to zero as $\lambda \nearrow \infty$. We now focus on the first term, $\frac{1}{\lambda} \log \mathcal{P}_{\underline{\lambda}}(\lambda)$. Since $\gamma m_1 < \epsilon < 1$, for all $i \leq \underline{\lambda}$,

$$\frac{i}{\lambda} \frac{\gamma m_1}{\epsilon} < \frac{\underline{\lambda}}{\lambda} \frac{\gamma m_1}{\epsilon} \leq \left(1 + \frac{1}{\lambda^{0.25}} \right) \frac{\gamma m_1}{\epsilon} \leq 1,$$

for large values of λ . Thus, we can write using Bernoulli's inequality that

$$\mathcal{P}_{\underline{\lambda}}(\lambda) \geq \prod_{i=1}^{\underline{\lambda}} \left(1 - i \frac{\gamma m_1}{\lambda} \right) = \prod_{i=1}^{\underline{\lambda}} \left(1 - \epsilon \frac{i \gamma m_1}{\epsilon \lambda} \right) \geq \prod_{i=1}^{\underline{\lambda}} (1 - \epsilon)^{\frac{i \gamma m_1}{\epsilon \lambda}}$$

for large values of λ . Therefore, by combining the trivial bound (A.15) and the above conclusions,

$$\delta \geq \max \left(-1, \frac{\gamma m_1}{2} \left[\frac{\log (1 - \epsilon)}{\epsilon} \right] \right) \rightarrow 0 \text{ as } \gamma m_1 \searrow 0.$$

Upper bounds: We have a complete proof of the large deviations of $\mathcal{P}(\lambda)$ for the case $\eta d > 2$. So it remains to prove the theorem for $0 < \eta d \leq 2$. We first prove the required upper bounds for the case $1 < \eta d \leq 2$. If $0 < a < 0.5$ and $\bar{\lambda} = \lceil \lambda(1 - \frac{1}{\lambda^a}) \rceil$, then from Lemma A.4, for any $\epsilon > 0$,

$$\mathcal{P}(\lambda) \leq N_{\lambda, \bar{\lambda}} \exp \left(-\frac{\gamma \bar{\lambda}^2 (m_1 - \epsilon)}{2 \lambda^{\eta d}} \right) [1 + o(1)]. \tag{A.16}$$

$$\tag{A.17}$$

Case: $\eta d > 1$. By applying log on both sides of (A.16) and then dividing by $\lambda^{2-\eta d}$, we see that

$$\frac{1}{\lambda^{2-\eta d}} \log \mathcal{P}(\lambda) \leq -\frac{\gamma (m_1 - \epsilon)}{2} \left(1 + \frac{1}{\lambda^a} \right)^2 + \frac{1}{\lambda^{2-\eta d}} \log N_{\lambda, \bar{\lambda}} + \frac{1}{\lambda^{2-\eta d}} \log [1 + o(1)].$$

As a consequence of Lemma A.3,

$$\limsup_{\lambda \rightarrow \infty} \frac{1}{\lambda^{2-\eta d}} \log \mathcal{P}(\lambda) \leq -\frac{\gamma (m_1 - \epsilon)}{2}.$$

Now take $\epsilon \searrow 0$.

In particular, consider the torus-hard-sphere model with $5/3 < \eta d \leq 2$. We can fix a such that $2 - \eta d < a < \frac{\eta d - 1}{2}$. From Lemma A.4, (A.16) holds with $\epsilon = 0$. Therefore,

$$\mathcal{P}(\lambda) \exp \left(\frac{\gamma m_1}{2} \lambda^{2-\eta d} \right) \leq N_{\lambda, \bar{\lambda}} \exp \left(O \left(\lambda^{2-\eta d - a} \right) \right) [1 + o(1)],$$

and hence

$$\limsup_{\lambda \rightarrow \infty} \left[\mathcal{P}(\lambda) \exp \left(\frac{\gamma m_1}{2} \lambda^{2-\eta d} \right) \right] \leq 1$$

from Lemma A.3.

Case: $0 < \eta d < 1$. Let $\underline{\lambda} = \lfloor \lambda^{\frac{1+\eta d}{2}} \rfloor$ and $N \sim Poi(\lambda)$. From the definition,

$$\mathcal{P}(\lambda) = \mathbb{E} [\mathcal{P}_N(\lambda)] \leq \mathbb{P} (N \leq \underline{\lambda}) + \mathbb{E} [\mathcal{P}_N(\lambda); N \geq \underline{\lambda} + 1]. \tag{A.18}$$

For any $\epsilon > 0$, let $H_n(\epsilon) := \left\{ \frac{1}{n} \sum_{i=1}^n R_i^d > \epsilon \right\}$. From (4.8),

$$\mathcal{P}_{n+1}(\lambda) \leq \mathbb{E} \left[\prod_{i=1}^n \left(1 - \frac{\gamma'}{\lambda^{\eta d}} \sum_{j=1}^i R_j^d \right)^+ \right] \leq \mathbb{P} \left(H_n^c \left(\frac{\lambda^{\eta d}}{\gamma' n} \right) \right) \leq \mathbb{P} \left(H_n^c \left(\frac{\lambda^{\eta d}}{\gamma' \underline{\lambda}} \right) \right),$$

where the second inequality holds because

$$\left(1 - \frac{\gamma'}{\lambda^{\eta d}} \sum_{j=1}^n R_j^d \right)^+ = 0$$

on $H_n \left(\frac{\lambda^{\eta d}}{\gamma' n} \right)$. Since $\eta d < (\eta d + 1)/2 < 1$, we see that $\frac{\lambda^{\eta d}}{\gamma' \underline{\lambda}} \searrow 0$ as $\lambda \nearrow \infty$, and thus for every $\epsilon > 0$ there exists λ_ϵ such that

$$\mathcal{P}_{n+1}(\lambda) \leq \mathbb{P} (H_n^c(\epsilon))$$

for all $\lambda > \lambda_\epsilon$ and $n > \underline{\lambda}$.

Suppose there is a constant $c > 0$ such that $R \geq c$. Then for all sufficiently small values of ϵ , $\mathbb{P} (H_n^c(\epsilon)) = 0$ for all $n > \underline{\lambda}$. Thus for large values of λ , $\mathcal{P}(\lambda) \leq \mathbb{P} (N \leq \underline{\lambda}) \leq e^{-\lambda} \lambda^{\underline{\lambda}}$, and from the definition of $\underline{\lambda}$,

$$\limsup_{\lambda \rightarrow \infty} \frac{1}{\lambda} \log \mathcal{P}(\lambda) \leq -1 + \limsup_{\lambda \rightarrow \infty} \left[\frac{\underline{\lambda} \log \lambda}{\lambda} \right] = -1.$$

So we can assume that $\mathbb{P}(R < \epsilon) > 0$ for every $\epsilon > 0$. Recall that $\mathbb{P}(R > 0) = 1$ and $\Lambda(\theta)$ is the logarithmic moment generating function of R^d . As a consequence of positivity of R , we see that $\Lambda(\theta) \searrow -\infty$ as $\theta \searrow -\infty$. Let $\Lambda^*(x) = \sup_{\theta \in \mathbb{R}} \{\theta x - \Lambda(\theta)\}$. As a consequence of the assumption that $\mathbb{P}(R < \epsilon) > 0$ for every $\epsilon > 0$, we can show that $\Lambda^*(x) \nearrow \infty$ as $x \searrow 0$. From Theorem 2.2.3 of [10],

$$\mathbb{P} (H_n^c(\epsilon)) \leq 2 \exp \left(-n \inf_{x \leq \epsilon} \Lambda^*(x) \right) = 2 \exp (-n \Lambda^*(\epsilon))$$

for all $n > \underline{\lambda}$ and $\epsilon < \mathbb{E}[R_1^d]$, where the last inequality holds because $\Lambda^*(x)$ is non-decreasing over $0 < x \leq \mathbb{E}[R_1^d]$. By (A.18),

$$\begin{aligned} \mathcal{P}(\lambda) &\leq \mathbb{P} (N \leq \underline{\lambda}) + \mathbb{P} (H_N^c(\epsilon); N \geq \underline{\lambda} + 1) \leq \mathbb{P} (N \leq \underline{\lambda}) + \mathbb{P} (H_N^c(\epsilon)) \\ &\leq \mathbb{P} (N \leq \underline{\lambda}) + 2 \exp \left(-\lambda \left(1 - e^{-\Lambda^*(\epsilon)} \right) \right) \end{aligned}$$

for all $\lambda \geq \lambda_\epsilon$. To conclude that $\limsup_{\lambda \rightarrow \infty} \frac{1}{\lambda} \log \mathcal{P}(\lambda) \leq -1$, we see from the definition of Poisson distribution and $\underline{\lambda}$ that

$$\mathbb{P}(N \leq \underline{\lambda}) = \sum_{n=0}^{\underline{\lambda}} e^{-\lambda} \frac{\lambda^n}{n!} \leq \underline{\lambda} e^{-\lambda} \left(\frac{\lambda^{\underline{\lambda}}}{\underline{\lambda}!} \right) \leq e^{-\lambda} \lambda^{\underline{\lambda}},$$

where we used the fact that $\lambda^{n-1}/(n-1)! < \lambda^n/n!$ for all $n < \lambda$. Hence,

$$\begin{aligned} \mathcal{P}(\lambda) &\leq 2 \exp\left(-\lambda \left(1 - e^{-\Lambda^*(\epsilon)}\right)\right) \left(1 + \lambda^{\underline{\lambda}} \exp\left(-\lambda \Lambda^*(\epsilon)\right)\right) \\ &= 2 \exp\left(-\lambda \left(1 - e^{-\Lambda^*(\epsilon)}\right)\right) \left(1 + \exp\left(-\lambda \left(\Lambda^*(\epsilon) - \frac{\underline{\lambda}}{\lambda} \log \lambda\right)\right)\right). \end{aligned}$$

From the definition of $\underline{\lambda}$, we see that $\frac{\underline{\lambda}}{\lambda} \log \lambda \searrow 0$ as $\lambda \nearrow \infty$. As a consequence, as $\lambda \nearrow \infty$, $\exp\left(-\lambda \left(\Lambda^*(\epsilon) - \frac{\underline{\lambda}}{\lambda} \log \lambda\right)\right)$ goes to zero. Therefore,

$$\limsup_{\lambda \rightarrow \infty} \frac{1}{\lambda} \log \mathcal{P}(\lambda) \leq -\left(1 - e^{-\Lambda^*(\epsilon)}\right).$$

We have the required result by taking $\epsilon \searrow 0$.

Case: $\eta d = 1$. It remains to show that

$$\delta \leq -\frac{1}{2} \left(1 - \frac{1}{\gamma' \bar{r}^d}\right)^2$$

if $R \equiv \bar{r}$ and $\gamma' \bar{r}^d > 1$. Since, from (4.8),

$$\mathcal{P}_{n+1}(\lambda) \leq \prod_{i=1}^n \left(1 - \frac{\gamma'}{\lambda} i \bar{r}^d\right)^+ = 0$$

for all $n > \lambda \frac{\lambda}{\gamma' \bar{r}^d}$, we have $\mathcal{P}(\lambda) \leq \mathbb{P}\left(N \leq \lambda \frac{\lambda}{\gamma' \bar{r}^d}\right)$. Now the proof is complete by Lemma A.1. □

A.2. Proof of Proposition 3.1

From [29], the intensity of the torus-hard-sphere model is given by

$$\rho(\lambda) = \frac{\sum_{n=1}^{\infty} n \frac{\lambda^n}{n!} \mathcal{P}_n(\lambda)}{\sum_{n \in \mathbb{N}_0} \frac{\lambda^n}{n!} \mathcal{P}_n(\lambda)} = \lambda \frac{\sum_{n \in \mathbb{N}_0} \frac{\lambda^n}{n!} \mathcal{P}_{n+1}(\lambda)}{\sum_{n \in \mathbb{N}_0} \frac{\lambda^n}{n!} \mathcal{P}_n(\lambda)},$$

where $\mathcal{P}_n(\lambda)$ is the non-overlapping probability of n uniformly and independently generated spheres with radius \bar{r}/λ^η .

Case $\eta d > 1$. In this regime, we show that $\rho(\lambda)$ is of the order of $\gamma' \bar{r}^d \lambda^{1-\eta d}$. Using inequalities (4.8) and (A.3),

$$\left(1 - n\gamma' \bar{r}^d \lambda^{-\eta d}\right) \mathcal{P}_n(\lambda) \geq \mathcal{P}_{n+1}(\lambda) \geq \left(1 - n\gamma' 4\bar{r}^d \lambda^{-\eta d}\right) \mathcal{P}_n(\lambda).$$

Therefore,

$$\rho(\lambda) \geq \lambda \frac{\sum_{n \in \mathbb{N}_0} \frac{\lambda^n}{n!} (1 - n4\bar{r}^d \lambda^{-\eta d}) \mathcal{P}_n(\lambda)}{\sum_{n \in \mathbb{N}_0} \frac{\lambda^n}{n!} \mathcal{P}_n(\lambda)} \geq \lambda - \gamma 4\bar{r}^d \lambda^{1-\eta d} \rho(\lambda),$$

and

$$\rho(\lambda) \leq \lambda \frac{\sum_{n \in \mathbb{N}_0} \frac{\lambda^n}{n!} (1 - n\bar{r}^d \lambda^{-\eta d}) \mathcal{P}_n(\lambda)}{\sum_{n \in \mathbb{N}_0} \frac{\lambda^n}{n!} \mathcal{P}_n(\lambda)} \leq \lambda - \gamma \bar{r}^d \lambda^{1-\eta d} \rho(\lambda).$$

Consequently,

$$\left(\frac{1}{1 + \gamma 4\bar{r}^d \lambda^{1-\eta d}} \right) \gamma \bar{r}^d \lambda^{1-\eta d} \leq \rho(\lambda) \gamma \bar{r}^d \lambda^{-\eta d} \leq \left(\frac{1}{1 + \gamma \bar{r}^d \lambda^{1-\eta d}} \right) \gamma \bar{r}^d \lambda^{1-\eta d},$$

and thus

$$\lim_{\lambda \nearrow \infty} \frac{\mathbf{VF}(\lambda)}{\gamma \bar{r}^d \lambda^{1-\eta d}} = 1.$$

Case $\eta d \leq 1$. We now show that $\lim_{\lambda \nearrow \infty} \mathbf{VF}(\lambda) \leq \rho^{\max} \gamma$ with equality if and only if $\eta d < 1$. Towards this end, we first consider another torus-hard-sphere model on $[0, \lambda^\eta / \bar{r}]^d$ with unit-radius spheres and absolutely continuous with respect to a κ -homogeneous PPP for some intensity $\kappa > 0$. Let $\rho(\kappa, \lambda)$ be the intensity of this new hard-sphere model. We can easily see that when $\kappa = \bar{r}^d \lambda^{1-\eta d}$, the fraction of the volume occupied by the spheres in the new hard-sphere model is also $\mathbf{VF}(\lambda)$.

The proofs of Propositions 1 and 2 in [29] can easily be modified to show that $\rho(\kappa, \lambda)$ is strictly increasing in κ for any fixed $\lambda > 0$, and

$$\lim_{\kappa \rightarrow \infty} \rho(\kappa, \lambda) = \rho^{\max},$$

where ρ^{\max} is the closest packing density. On the other hand, by fixing κ , we can further show, using [29], that the limit $\lim_{\lambda \rightarrow \infty} \rho(\kappa, \lambda)$ exists and is equal to the intensity of the stationary hard-sphere model on \mathbb{R}^d with unit-radius spheres and the reference PPP being κ -homogeneous. (In fact, the difference between $\rho(\kappa, \lambda)$ and the limit $\lim_{\lambda \rightarrow \infty} \rho(\kappa, \lambda)$ is known to be insignificant for large values of λ ; see, for example, [5].)

From the above discussion, when $\eta d < 1$, for sufficiently small $\epsilon > 0$, there exist constants κ_ϵ and λ_ϵ such that $\rho(\kappa, \lambda) > \rho^{\max} - \epsilon$ for all $\lambda > \lambda_\epsilon$ and $\kappa > \kappa_\epsilon$. If we take $\kappa = \bar{r}^d \lambda^{1-\eta d}$, since $\eta d < 1$,

$$\lim_{\lambda \rightarrow \infty} \mathbf{VF}(\lambda) = \lim_{\lambda \rightarrow \infty} \left[\rho(\bar{r}^d \lambda^{1-\eta d}, \lambda) \gamma \right] = \rho^{\max} \gamma,$$

which is the maximum packing intensity.

Finally, if $\eta d = 1$ and $\kappa = \bar{r}^d$, the limit $\lim_{\lambda \rightarrow \infty} \rho(\bar{r}^d, \lambda)$ is strictly less than ρ^{\max} . Hence, we have $\lim_{\lambda \rightarrow \infty} \mathbf{VF}(\lambda) < \rho^{\max} \gamma$.

A.3. Proof of Proposition 4.1

Let N' be the number of spheres generated sequentially, independently and identically, before seeing an overlap. Let $N \sim \text{Poi}(\lambda)$ independently of N' . Then from the construction of Algorithm 4.1,

$$c \mathbb{E} \left[\sum_{n=1}^{\min(N, N')} (n-1) \right] \leq \text{Citr}(\lambda) \leq c' \left(\log(\lambda) + \mathbb{E} \left[\sum_{n=1}^{\min(N, N')} (n-1) \right] \right) \tag{A.19}$$

for some positive constants c and c' . In the above expression, $\log(\lambda)$ appears because the cost to generate a sample of N is at most of order $\log(\lambda)$ (see e.g. [11]). Observe that

$$\mathbb{E} \left[\sum_{n=1}^{\min(N, N')} (n-1) \right] = \mathbb{E} \left[\sum_{n=0}^{N-1} nI(N' > n) \right] = \mathbb{E} \left[\sum_{n=1}^{N-1} n\mathcal{P}_n(\lambda) \right], \tag{A.20}$$

where the last equality follows from the fact that $\mathbb{P}(N' > n) = \mathcal{P}_n(\lambda)$.

Upper bound: For $\eta d \geq 2$, since $\mathcal{P}_n(\lambda) \leq 1$, we can bound (A.20) from above by a constant times $\mathbb{E}[N^2]$, which is further bounded from above by a constant times λ^2 . Therefore we just need to consider the case $\eta d < 2$. From (A.2),

$$\mathcal{P}_n(\lambda) = \mathbb{E}_{\tilde{\mu}_n} \left[\prod_{i=1}^n (1 - B_i) \right].$$

As a consequence of (4.8),

$$\mathcal{P}_n(\lambda) \leq \mathbb{E} \left[\exp \left(-\frac{\gamma'}{\lambda \eta d} \sum_{1 \leq j < i \leq n} R_j^d \right) \right] = \mathbb{E} \left[\exp \left(-\frac{\gamma' \bar{r}^d}{\lambda \eta d} \sum_{1 \leq j < i \leq n} \frac{R_j^d}{\bar{r}^d} \right) \right].$$

Let $\alpha = \mathbb{E}[R_1^d]$. Since \bar{r} is an upper bound on the R_i , by Hoeffding’s inequality (Lemma A.2) on the sequence $\{R_1^d/\bar{r}^d, \dots, R_n^d/\bar{r}^d\}$ with $\epsilon = \alpha/2\bar{r}^d$, $k = 2$, and $g(x, y) = x$, we have

$$\begin{aligned} \mathcal{P}_n(\lambda) &\leq \exp \left(-\frac{\gamma'}{2\lambda \eta d} \frac{n(n-1)}{2} \alpha \right) + \mathbb{P} \left(\sum_{1 \leq j < i \leq n} \frac{R_j^d}{\bar{r}^d} \leq \frac{\alpha}{2\bar{r}^d} \right) \\ &\leq \exp \left(-\frac{\gamma'(n-1)^2}{4\lambda \eta d} \alpha \right) + \exp \left(-\frac{n\alpha^2}{4\bar{r}^{2d}} \right). \end{aligned}$$

Let $a = \sqrt{\frac{2\lambda \eta d}{\gamma' \alpha}}$. Then from the above expression,

$$\sum_{n=1}^{\infty} n \mathcal{P}_n(\lambda) \leq \sum_{n=1}^{\infty} n \exp \left(-\frac{(n-1)^2}{2a^2} \right) + \sum_{n=1}^{\infty} n \exp \left(-\frac{n\alpha^2}{4\bar{r}^{2d}} \right). \tag{A.21}$$

Select λ large enough so that $b > 0$. Then with $p = 1 - \exp(-\alpha^2/(4\bar{r}^{2d}))$, the second term on the right side of (A.21) is $1/p$ times $\mathbb{E}[Z]$ for a geometric random variable Z with success probability p and support $\{1, 2, 3, \dots\}$. Since $\mathbb{E}[Z] = 1/p$, the term $\sum_{n=1}^{\infty} n \exp(-n\alpha^2/(4\bar{r}^{2d}))$ is bounded from above by a constant.

On the other hand, since

$$n \exp \left(-(n-1)^2/(2a^2) \right) \leq \int_{n-2}^{n-1} (x+2) \exp \left(-\frac{x^2}{2a^2} \right)$$

for any non-negative integer n , we can write that

$$\sum_{n=1}^{\infty} n \exp \left(-\frac{(n-1)^2}{2a^2} \right) \leq 1 + \int_0^{\infty} (x+2) \exp \left(-\frac{x^2}{2a^2} \right) dx = 1 + a\sqrt{\frac{\pi}{2}} \mathbb{E}[(|Z|+2)],$$

where Z is a Gaussian random variable with mean 0 and variance a^2 . Since $\mathbb{E}[|Z|] = a\sqrt{2/\pi}$, using the definition of a , the first term in (A.21) is bounded from above by a constant times $\lambda^{\eta d}$. Thus, the required upper bound on (A.19) is established as a consequence of (A.20) and (A.21).

Lower bound: Let $\epsilon' = \min(1, \eta d/2)$. Then from (A.20),

$$C_{\text{itr}}(\lambda) \geq \mathbb{P}\left(N > \lambda^{\epsilon'}/2\right) \sum_{n=1}^{\lfloor \lambda^{\epsilon'}/2 \rfloor + 1} n \mathcal{P}_n(\lambda) \geq c \mathcal{P}_{\lfloor \lambda^{\epsilon'}/2 \rfloor}(\lambda) \left(\frac{\lambda^{\epsilon'}}{2}\right)^2$$

for a constant $c > 0$. From (A.4),

$$\mathcal{P}_{\lfloor \lambda^{\epsilon'}/2 \rfloor}(\lambda) \geq \exp\left(-\frac{\lambda^{\epsilon'}}{2} \sum_{j=1}^{\infty} \left(\frac{\gamma \lambda^{\epsilon'}}{2\lambda^{\eta d}}\right)^j \frac{m_j}{j(j+1)}\right) \geq \exp\left(-\frac{\lambda^{\epsilon'}}{2} \sum_{j=1}^{\infty} \left(\frac{\gamma \lambda^{\epsilon'}}{2\lambda^{\eta d}}\right)^j \frac{m_j}{j}\right),$$

where $m_j = \mathbb{E}[(R_1 + R_2)^{jd}]$. Note that $m_j \leq (2\bar{r})^{jd}$ since $R \leq \bar{r}$. Therefore,

$$\mathcal{P}_{\lfloor \lambda^{\epsilon'}/2 \rfloor}(\lambda) \geq \exp\left(-\frac{\lambda^{\epsilon'}}{2} \sum_{j=1}^{\infty} \frac{1}{j} \left(\frac{\gamma(2\bar{r}^d)}{2\lambda^{\eta d - \epsilon'}}\right)^j\right).$$

Using the Taylor expansion of $\log(1 - x)$ for $0 < x < 1$, along with the fact that

$$\frac{\gamma(2\bar{r}^d)}{2\lambda^{\eta d - \epsilon'}} < 1$$

for sufficiently large values of λ , we have

$$\mathcal{P}_{\lfloor \lambda^{\epsilon'}/2 \rfloor}(\lambda) \geq \exp\left(\frac{\lambda^{\epsilon'}}{2} \log\left(1 - \frac{\gamma(2\bar{r}^d)}{2\lambda^{\eta d - \epsilon'}}\right)\right) = \left(1 - \frac{\gamma(2\bar{r}^d)}{2\lambda^{\eta d - \epsilon'}}\right)^{\frac{\lambda^{\epsilon'}}{2}}.$$

From the definition of ϵ' ,

$$\lim_{\lambda \rightarrow \infty} \left[\left(1 - \frac{\gamma(2\bar{r}^d)}{2\lambda^{\eta d - \epsilon'}}\right)^{\frac{\lambda^{\epsilon'}}{2}} \right] = \begin{cases} 1 & \text{if } \eta d > 2, \\ \exp(-\gamma(2\bar{r})^d/4) & \text{if } 0 < \eta d \leq 2. \end{cases}$$

In addition, from Lemma A.1, $\lim_{\lambda \rightarrow \infty} \mathbb{P}(N > \lambda^{\epsilon'}/2) = 1$. Therefore, there exists a constant c such that $C_{\text{itr}}(\lambda) \geq c \lambda^{2\epsilon'}$. The proof of Proposition 4.1 is then complete by Theorem 3.1 and (4.1).

A.4. Proof of Proposition 4.3

First note that the sphere volume is at most a constant times the cell volume for all λ . Thus, after generating a sphere, the complexity of relabeling cells around the center of the new sphere plus the complexity of the overlap check is a constant. For $\eta d > 1$, the number of

spheres generated in an iteration of Algorithm 4.5 is stochastically dominated by a Poisson random variable with rate λ . Therefore, there exists a constant c such that $\tilde{C}_{\text{itr}}(\lambda) \leq c\lambda$. On the other hand, if $0 < \eta d \leq 1$, the expected number of spheres generated per iteration is of order $\lambda^{\eta d}$, because the expected volume of each sphere is of order $1/\lambda^{\eta d}$. It is clear that there exists a constant $c > 0$ such that $\tilde{C}_{\text{itr}}(\lambda) \leq c\lambda^{\eta d}$. Thus, by (4.6) and Proposition 4.1,

$$\mathcal{T}_{\text{ISAR}} \leq c \frac{\lambda^{\min(1, \eta d)}}{P_{\text{acc}}(\lambda)} = c \mathbb{E}[\tilde{\sigma}(N)] \frac{\lambda^{\min(1, \eta d)}}{\mathcal{P}(\lambda)}.$$

Thus, (4.11) holds, because $\tilde{\sigma}(n) = \delta_n$ for each $n \in \mathbb{N}_0$. Furthermore, from the definition of $\tilde{\sigma}(\cdot)$ and N ,

$$\mathbb{E}[\delta_N] \leq \mathbb{E} \left[\exp \left(- \sum_{i=0}^N (i-1) \gamma' \frac{\bar{r}^d}{\lambda^{\eta d}} \right) \right] = \mathbb{E} \left[\exp \left(- \gamma' \frac{\bar{r}^d}{2\lambda^{\eta d}} \frac{(N-1)N}{2} \right) \right].$$

By the Chernoff bound (Lemma A.1), for any $0 < \epsilon < 1$,

$$\begin{aligned} \mathbb{E} \left[\exp \left(- \gamma' \frac{\bar{r}^d}{\lambda^{\eta d}} \frac{(N-1)N}{2} \right) \right] &\leq \mathbb{E} \left[\exp \left(- \gamma' \frac{\bar{r}^d}{\lambda^{\eta d}} \frac{(N-1)N}{2} \right); N > \lambda(1-\epsilon) \right] \\ &\quad + \mathbb{P}(N \leq \lambda(1-\epsilon)) \\ &\leq \exp \left(- \frac{\gamma' \bar{r}^d}{2} (1-\epsilon) \lambda^{2-\eta d} \right) + \exp \left(- \frac{\lambda \epsilon^2}{2} \right). \end{aligned}$$

If $\eta d > 1$, then the second term on the right-hand side of the above expression decreases faster than the first term, and thus the claim holds true. For $\eta d = 1$, take $\epsilon = 1/2$; then we have the required result with $b = \min(1/8, \gamma \bar{r}^d/4)$. Furthermore, if $0 < \eta d < 1$ then the first term decreases faster than the second one, and hence the proof is completed by taking $b = 1/2$.

A.5. Proof of Theorem 5.1

To derive the lower bound on \mathcal{T}_{DC1} , we view the entire dominating process \mathbf{D} as a Poisson Boolean model on a higher-dimensional space and use an extension of the FKG inequality [31] (alternatively, see Theorem 2.2 in [31]). Let $s_0 = 0$, and let s_i be the instant of the i th arrival in the dominating process after time zero. Let $C(\mathbf{x}, \mathbf{x}^u, \mathbf{x}^l)$ be the running time complexity of updating the dominating, upper bound, and lower bound processes at the instant of an arrival when their respective states are \mathbf{x} , \mathbf{x}^u , and \mathbf{x}^l .

Since $U_0(0) = D(0)$ and $L_0(0) = \emptyset$, on $\bigcap_{j=0}^i \{D(s_j) \notin \mathcal{A}\}$, for all $t \leq s_i$,

$$L_0(t) = \emptyset \quad \text{and} \quad U_0(t) = D(t). \tag{A.22}$$

Thus, $L_0(t) \neq U_0(t)$ for all $t \leq s_i$ on $\bigcap_{j=0}^i \{D(s_j) \notin \mathcal{A}\}$. Now take

$$\tau = \inf \{i \in \mathbb{N}_0 : D(s_i) \in \mathcal{A}\}.$$

From the above conclusion, it is clear that $N^f \geq \tau$. Then,

$$\begin{aligned} \mathcal{T}_{DC1} &\geq \mathbb{E} \left[\sum_{i=0}^{N^f} C(D(s_i), U_0(s_i), L_0(s_i)) \right] \\ &\geq \mathbb{E} \left[\sum_{i=0}^{\tau} C(D(s_i), U_0(s_i), L_0(s_i)) \right] \\ &= \sum_{i=0}^{\infty} \mathbb{E} \left[C(D(s_i), U_0(s_i), L_0(s_i)); \tau \geq i \right] \\ &= \sum_{i=0}^{\infty} \mathbb{E} \left[C(D(s_i), U_0(s_i), L_0(s_i)); \bigcap_{j=0}^{i-1} \{D(s_j) \notin \mathcal{A}\} \right] \\ &= \sum_{i=0}^{\infty} \mathbb{E} \left[C(D(s_i), D(s_i), \emptyset); \bigcap_{j=0}^{i-1} \{D(s_j) \notin \mathcal{A}\} \right], \end{aligned}$$

where $I \left(\bigcap_{j=0}^{-1} \{D(s_j) \notin \mathcal{A}\} \right) = 1$ and the last equality follows from (A.22).

Suppose that \mathcal{D} is the state space of the entire process $\{D(t) : t \in \mathbb{R}\}$. Then we can define a simple partial order on \mathcal{D} as follows. For any $\omega, \omega' \in \mathcal{D}$, we say $\omega \leq \omega'$ if and only if every sphere present in ω is also present in ω' ; that is, either $\omega' = \omega$ or ω' is obtained by adding spheres to ω . We then define the following notion of increasing functions: a real-valued function on \mathcal{D} is increasing if $f(\omega) \leq f(\omega')$ for all $\omega, \omega' \in \mathcal{D}$ such that $\omega \leq \omega'$.

At each arrival, if there are n points in the upper bounding process, the cost of deciding whether to accept the new point is at least the cost of checking the overlap condition in the upper bounding process, and that cost is of order n . Therefore, $C(D(s_i), U_0(s_i), \emptyset) = c|U_0(s_i)|$ for some constant c . Since $|U_0(s_i)|$ is a non-decreasing function on \mathcal{D} as per the partial order stated above, by the FKG inequality (Theorem 2.2 in [31]),

$$\mathbb{E} \left[C(D(s_i), D(s_i), \emptyset); \bigcap_{j=0}^{i-1} \{D(s_j) \notin \mathcal{A}\} \right]$$

is bounded from below by

$$\mathbb{E} \left[C(D(s_i), D(s_i), \emptyset) \right] \prod_{j=0}^{i-1} \mathbb{P} (D(s_j) \notin \mathcal{A}) .$$

Thus,

$$\mathcal{T}_{DC1} \geq \frac{\mathbb{E} \left[C(D(s_1), D(s_1), \emptyset) \right]}{1 - \mathbb{P} (D(s_1) \notin \mathcal{A})} = \frac{\mathbb{E}_{\mu^0} [|\mathbf{X}|]}{\mathbb{P}_{\mu^0} (\mathbf{X} \in \mathcal{A})} = c \frac{\lambda}{\mathcal{P}(\lambda)}$$

for some constant $c > 0$. Then (5.2) follows from (4.1) and (4.2). The proof is completed using Theorem 3.1.

References

- [1] ADAMS, D. J. (1974). Chemical potential of hard-sphere fluids by Monte Carlo methods. *Molec. Phys.* **28**, 1241–1252.
- [2] AMORÓS, J. AND RAVI, S. (2013). On the application of the Carnahan–Starling method for hard hyperspheres in several dimensions. *Phys. Lett. A* **377**, 2089–2092.
- [3] ASMUSSEN, S. AND GLYNN, P. W. (2007). *Stochastic Simulation: Algorithms and Analysis* Springer, New York.
- [4] AURENHAMMER, F. (1987). Power diagrams: properties, algorithms and applications. *SIAM J. Comput.* **16**, 78–96.
- [5] BADDELEY, A. AND NAIR, G. (2012). Fast approximation of the intensity of Gibbs point processes. *Electron. J. Statist.* **6**, 1155–1169.
- [6] BARANAU, V. AND TALLAREK, U. (2016). Chemical potential and entropy in monodisperse and polydisperse hard-sphere fluids using Widom’s particle insertion method and a pore size distribution-based insertion probability. *J. Chem. Phys.* **144**, 214503.
- [7] BERTHELSEN, K. K. AND MØLLER, J. (2006). Bayesian analysis of Markov point processes. In *Case Studies in Spatial Point Process Modeling*, Springer, New York, pp. 85–97.
- [8] BERTHELSEN, K. K. AND MØLLER, J. (2008). Non-parametric Bayesian inference for inhomogeneous Markov point processes. *Austral. N. Z. J. Statist.* **50**, 257–272.
- [9] BOUTE, R. T. (1992). The Euclidean definition of the functions div and mod . *ACM Trans. Program. Lang. Systems* **14**, 127–144.
- [10] DEMBO, A. AND ZEITOUNI, O. (2010). *Large Deviations Techniques and Applications*. Springer, Berlin.
- [11] DEVROYE, L. (1986). *Nonuniform Random Variate Generation*. Springer, New York.
- [12] FERRARI, P. A., FERNÁNDEZ, R. AND GARCIA, N. L. (2002). Perfect simulation for interacting point processes, loss networks and Ising models. *Stoch. Process. Appl.* **102**, 63–88.
- [13] FILL, J. A. (1998). An interruptible algorithm for perfect sampling via Markov chains. *Ann. Appl. Prob.* **8**, 131–162.
- [14] FILL, J. A., MACHIDA, M., MURDOCH, D. J. AND ROSENTHAL, J. S. (2000). Extension of Fill’s perfect rejection sampling algorithm to general chains. *Random Structures Algorithms* **17**, 290–316.
- [15] FOSS, S., JUNEJA, S., MANDJES, M. AND MOKA, S. (2015). Spatial loss systems: exact simulation and rare event behavior. *ACM SIGMETRICS Performance Evaluation Rev.* **43**, 3–6.
- [16] GARCIA, N. L. (2000). Perfect simulation of spatial processes. *Resenhas* **4**, 283–325.
- [17] GRIMMETT, G. (1999). *Percolation*, 2nd edn. Springer, Berlin.
- [18] HAENGGI, M. (2012). *Stochastic Geometry for Wireless Networks*. Cambridge University Press.
- [19] HÄGGSTRÖM, O., VAN LIESHOUT, M.-C. N. AND MØLLER, J. (1999). Characterization results and Markov chain Monte Carlo algorithms including exact simulation for some spatial point processes. *Bernoulli* **5**, 641–658.
- [20] HIRSCH, C., MOKA, S. B., TAIMRE, T. AND KROESE, D. P. (2021). Rare events in random geometric graphs. *Methodology Comput. Appl. Prob.*, doi: [10.1007/s11009-021-09857-7](https://doi.org/10.1007/s11009-021-09857-7).
- [21] HOEFFDING, W. (1963). Probability inequalities for sums of bounded random variables. *J. Amer. Statist. Assoc.* **58**, 13–30.
- [22] HUBER, M. L. (2012). Spatial birth–death swap chains. *Bernoulli* **18**, 1031–1041.
- [23] HUBER, M. L. (2016). *Perfect Simulation*. CRC Press, Boca Raton, FL.
- [24] JUNEJA, S. AND MANDJES, M. (2013). Overlap problems on the circle. *Adv. Appl. Prob.* **45**, 773–790.
- [25] KENDALL, W. S. (1998). Perfect simulation for the area-interaction point process. In *Probability towards 2000 (New York, 1995)* (Lecture Notes Statist., Vol. 128), Springer, New York, pp. 218–234.
- [26] KENDALL, W. S. (2015). Introduction to coupling-from-the-past using R . In *Stochastic Geometry, Spatial Statistics and Random Fields* (Lecture Notes Math., Vol. 2120), Springer, Cham, pp. 405–439.
- [27] KENDALL, W. S. AND MØLLER, J. (2000). Perfect simulation using dominating processes on ordered spaces, with application to locally stable point processes. *Adv. Appl. Prob.* **32**, 844–865.
- [28] LIEB, E. AND MATTIS, D. (1966). *Mathematical Physics in One Dimension: Exactly Soluble Models of Interacting Particles*. Academic Press, New York.
- [29] MASE, S. et al. (2001). Packing densities and simulated tempering for hard core Gibbs point processes. *Ann. Inst. Statist. Math.* **53**, 661–680.
- [30] MAYER, J. AND MAYER, M. (1940). *Statistical Mechanics*. John Wiley, New York.
- [31] MEESTER, R. AND ROY, R. (1996). *Continuum Percolation*. Cambridge University Press.
- [32] MØLLER, J., HUBER, M. L. AND WOLPERT, R. L. (2010). Perfect simulation and moment properties for the Matérn type III process. *Stoch. Process. Appl.* **120**, 2142–2158.
- [33] MØLLER, J. AND HELISOVÁ, K. (2008). Power diagrams and interaction processes for unions of discs. *Adv. Appl. Prob.* **40**, 321–347.

- [34] MØLLER, J., PETTITT, A. N., REEVES, R. AND BERTHELSEN, K. K. (2006). An efficient Markov chain Monte Carlo method for distributions with intractable normalising constants. *Biometrika* **93**, 451–458.
- [35] PATHRIA, R. (1972). *Statistical Mechanics*. Elsevier, Amsterdam.
- [36] PROPP, J. G. AND WILSON, D. B. (1996). Exact sampling with coupled Markov chains and applications to statistical mechanics. *Random Structures Algorithms* **9**, 223–252.
- [37] SALSBERG, Z. W., ZWANZIG, R. W. AND KIRKWOOD, J. G. (1953). Molecular distribution functions in a one-dimensional fluid. *J. Chem. Phys.* **21**, 1098–1107.
- [38] SENGER, B., VOEGEL, J.-C. AND SCHAAF, P. (2000). Irreversible adsorption of colloidal particles on solid substrates. *Colloids Surfaces A* **165**, 255–285.
- [39] STOYAN, D., KENDALL, W. S. AND MECKE, J. (1987). *Stochastic Geometry and Its Applications*. John Wiley, Chichester.
- [40] TARJUS, G., SCHAAF, P. AND TALBOT, J. (1990). Generalized random sequential adsorption. *J. Chem. Phys.* **93**, 8352–8360.
- [41] THÖNNES, E. (1999). Perfect simulation of some point processes for the impatient user. *Adv. Appl. Prob.* **31**, 69–87.