Estimating hitting probabilities of an interacting particle system on a graph

T. Brereton, [a], **C. Eckard,** [a], **V. Schmidt,** [a]

[a] *Institute of Stochastics, University of Ulm, PLZ 89069, Ulm Germany*
Email: timothy.brereton@uni-ulm.de

Abstract:

We consider the problem of estimating hitting probabilities related to a class of interacting particle systems. These systems, in which two types of particles — 'electrons' and 'holes' — move on a graph, are simplified versions of models describing charge transport in disordered materials. The probability of interest is the probability that an electron reaches a certain region of the graph before colliding with a hole. We provide a detailed description of our model and explain how it can be simulated. Next, we give a brief introduction to importance sampling, which we use to improve the efficiency of our estimators. To our knowledge, importance sampling has not yet been used to estimate probabilities related to interacting particle systems. We describe how importance sampling can be used to improve the efficiency of estimators of hitting time probabilities involving discrete time Markov chains. We then use importance sampling to estimate the probability that we are interested in. In doing so, we observe that there are a number of complexities that arise when working with interacting particle systems. We describe some simple heuristics for implementing importance sampling. These heuristics make minimal changes to the probability measure under which the original system evolves. We consider a specific example of our problem and investigate the effectiveness of the importance sampling approach. We show that our estimators outperform standard Monte Carlo estimators. Finally, we describe possible future work, which includes a more sophisticated importance sampling approach that uses 'locally optimal' changes of measure.

Keywords: *Importance sampling, Monte Carlo, interacting particle systems, hitting time, rare event simulation, Markov chains.*

1 INTRODUCTION

In this paper, we apply importance sampling to an interacting particle system on a graph. To the best of our knowledge, such systems have not yet been considered in the rare event simulation literature. Interacting particle systems consist of a number of randomly moving particles which interact with one another, usually according to simple rules. A hallmark of such systems is that they often exhibit complex and interesting behavior on a macroscopic level. Interacting particle systems have applications in many fields and, in particular, can be used to model a number of fundamental physical processes. From a computational perspective, the most significant challenge when dealing with interacting particle systems is that they have extremely large state spaces. For a general introduction to the topic, see Liggett (2005).

The specific problem we consider has its origins in the physics of organic solar cells. Often, charge transport in such materials is modeled using a continuous time Markov chain, which describes the locations of charge carriers as they 'hop' on the vertices of a geometric graph; see Brereton et al. (2014) and Stenzel et al. (2014). One of the aims when designing efficient solar cells is to limit recombination, which causes a charge carrier to disappear before it reaches an electrode. In Markov chain models, an electron (one type of charge carrier) is treated as 'recombining' if it collides with a hole (another type of charge carrier). When studying recombination, a quantity of interest is the probability that an electron starting at a given location reaches an electrode without recombining. In general, such probabilities must be calculated using Monte Carlo methods. Because of the size and complexity of the systems being considered, the computational cost of the simulations is very large. Thus, it is of great value to develop tools that can reduce computational cost. One possible tool is importance sampling.

There are a number of challenges that must be addressed when applying importance sampling to interacting particle systems. We consider a simple problem that shares many features with realistic models of charge transport. We describe a technique by which importance sampling can be effectively applied to interacting particle systems and show that the resulting estimators outperform standard Monte Carlo. Finally, we briefly discuss some more sophisticated approaches that will be investigated in the future.

2 PROBLEM DESCRIPTION

The problem we consider can be described as follows. Two types of particles, 'electrons' and 'holes', move on a connected graph, $G = (V, E)$, according to a continuous time Markov chain. We consider the case where there is one electron and K holes, where $K > 0$ is an arbitrary integer. No more than one hole is allowed to occupy a vertex at any given time. When the electron collides with a hole, both the electron and the hole are destroyed. We wish to determine the probability that the electron is able to reach an 'electrode', $A \subset V$, before it is destroyed by a hole. The problem setting is illustrated in Figure 1.

More formally, the position of the electron is described by a V-valued continuous time Markov chain, $\{E_t\}_{t\geq 0}$, that jumps with rate 1 according to the transition matrix P_1. The positions of the holes are described by the V-valued continuous time stochastic processes, $\{H_t^{(1)}\}_{t\geq 0}, \ldots, \{H_t^{(K)}\}_{t\geq 0}$. Each of these processes attempts to jump at rate 1 according to the transition matrix P_2. However, if the vertex to which the process tries to jump is already occupied by a hole, the process remains where it is. Thus, the collection of processes, $\{H_t^{(1)}\}_{t\geq 0}, \ldots, \{H_t^{(K)}\}_{t\geq 0}$, forms an exclusion process (see Grimmett (2010), Liggett (2005)). Note that the complete system $\{S_t\}_{t\geq 0} = \{(E_t, H_t^{(1)}, \ldots, H_t^{(K)})\}_{t\geq 0}$ is a continuous time Markov chain taking values in a subset of V^{K+1}. The system starts in an initial configuration $(e_0, h_0^{(1)}, \ldots, h_0^{(K)})$.

We are interested in the event that $\{E_t\}_{t\geq 0}$ collides with one of the $\{H_t^{(k)}\}_{t\geq 0}$ before it reaches a target set $A \subset V$. More precisely, we define two hitting times, $\tau_{\mathsf{elec}} = \inf\{t \geq 0 : E_t \in A\}$ and $\tau_{\mathsf{recomb}} = \inf\left\{t \geq 0 : E_t = H_t^{(k)} \text{ for some } k \in \{1, \ldots, K\}\right\}$. The probability of interest is $\mathbb{P}(\tau_{\mathsf{elec}} < \tau_{\mathsf{recomb}})$. In order to determine this probability, it is sufficient to consider the discrete time Markov chain, $\{\widetilde{S}_n\}_{n\in\mathbb{N}_0} = \{(\widetilde{E}_n, \widetilde{H}_n^{(1)}, \ldots, \widetilde{H}_n^{(K)})\}_{n\in\mathbb{N}_0}$, which records the state of $\{S_t\}_{t\geq 0}$ after each attempted jump by a particle. In practice, $\{\widetilde{S}_n\}_{n\in\mathbb{N}_0}$ is updated at the $(n+1)$th step as follows.

Algorithm 2.1 (Generating $\widetilde{S}_{n+1}$ given $\widetilde{S}_n = (e_n, h_n^{(1)}, \ldots, h_n^{(K)})$)**.**

1. Draw $U \sim \mathcal{U}(0, 1)$.
2. If $U \leq (K+1)^{-1}$, move the electron. That is, draw $\widetilde{E}_{n+1} \sim P_1(e_n, \cdot)$ and set $\widetilde{H}_{n+1}^{(k)} = h_n^{(k)}$ for all

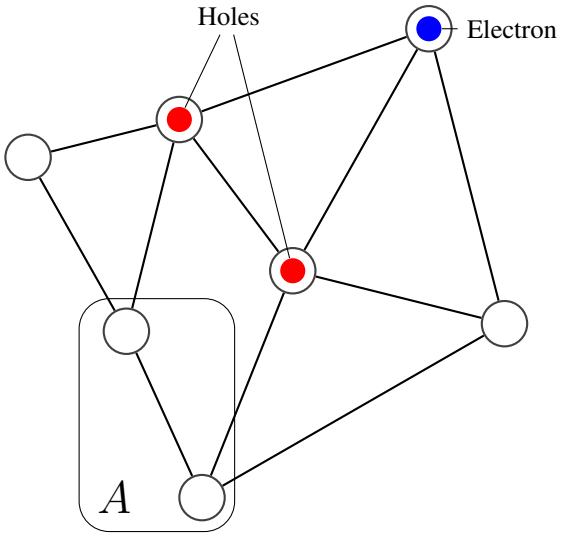

Figure 1. Illustration of the problem setting. The holes (colored red) and electron (colored blue) move about on the vertices of the graph. We wish to calculate the probability that the electron reaches the region A (the electrode) before it collides with a hole.

$k \in \{1, \dots, K\}$.

3. If $U > (K+1)^{-1}$, attempt to move a hole. That is, draw k^* uniformly from $\{1, \dots, K\}$. Draw the proposed move, $H^{\text{prop}} \sim P_2(h_n^{(k^*)}, \cdot)$.
   - If $H^{\text{prop}} \neq h_n^{(k)}$ for all $k \in \{1, \dots, K\}$, then set $\widetilde{H}_{n+1}^{(k^*)} = H^{\text{prop}}$, $\widetilde{E}_{n+1} = e_n$ and $\widetilde{H}_{n+1}^{(k)} = h_n^{(k)}$ for all $k \in \{1, \dots, K\} \setminus \{k^*\}$.
   - Otherwise, set $\widetilde{E}_{n+1} = e_n$ and $\widetilde{H}_{n+1}^{(k)} = h_n^{(k)}$ for all $k \in \{1, \dots, K\}$.

Because of the combinatorial complexity of the state space of $\{\widetilde{S}_n\}_{n \in \mathbb{N}_0}$, it is not practical to work with the transition matrix of $\{\widetilde{S}_n\}_{n \in \mathbb{N}_0}$ explicitly. However, observe that only three types of event can occur in one step: an electron can move, a hole can attempt to move and successfully move to another vertex, or a hole can attempt to move and remain where it is. We write $p^E(v \,|\, e, h^{(1)}, \dots, h^{(K)})$ for the conditional probability that the electron moves to $v \in V$ given the current state of the system is $(e, h^{(1)}, \dots, h^{(K)})$. Thus,

$$p^E(v \,|\, e, h^{(1)}, \dots, h^{(K)}) = \frac{1}{K+1} P_1(e, v).$$

We write $p^H(k, v \,|\, e, h^{(1)}, \dots, h^{(K)})$ for the conditional probability that the kth hole is selected to move and moves to $v \in V$ given the current state of the system is $(e, h^{(1)}, \dots, h^{(K)})$. Thus,

$$p^H(k, v \,|\, e, h^{(1)}, \dots, h^{(K)}) = \begin{cases} \frac{1}{K+1} \left[P_2(h^{(k)}, v) \, \mathbb{I}\left(v \neq h^{(i)} \text{ for all } i \in \{1, \dots, K\}\right) \right] & \text{if } v \neq h^{(k)}, \\ \frac{1}{K+1} \left[1 - \sum_{v' \in V \setminus \{v\}} P_2(h^{(k)}, v') \, \mathbb{I}\left(v' \neq h^{(i)} \text{ for all } i \in \{1, \dots, K\}\right) \right] & \text{if } v = h^{(k)}, \end{cases}$$

where $\mathbb{I}(A)$ denotes the indicator function of the set A. Note that $\sum_{v \in V} p^E(v \,|\, e, h^{(1)}, \dots, h^{(K)}) + \sum_{k=1}^{K} \sum_{v \in V} p^H(k, v \,|\, e, h^{(1)}, \dots, h^{(K)}) = 1$. Given the probabilities $p^E(v \,|\, e, h^{(1)}, \dots, h^{(K)})$ and $p^H(k, v \,|\, e, h^{(1)}, \dots, h^{(K)})$ for all $v \in V$ and $k \in \{1, \dots, K\}$, the dynamics of the Markov chain $\{\widetilde{S}_n\}_{n \in \mathbb{N}_0}$ are fully determined.

3 Rare event simulation

The problem described above is actually a hitting time problem. This is a class of problems that have been extensively studied in the rare event simulation literature. The classical setting is as follows. Let $\{X_n\}_{n \in \mathbb{N}_0}$

be a discrete time Markov chain with finite state space $\mathcal{X}$, transition matrix P and $X_0 = x_0$ for some $x_o \in \mathcal{X}$. We wish to estimate $\ell = \mathbb{P}(\tau_A < \tau_B)$, where A and B are disjoint sets in $\mathcal{X}$, $\tau_A = \inf\{n \in \mathbb{N}_0 : X_n \in A\}$, $\tau_B = \inf\{n \in \mathbb{N}_0 : X_n \in B\}$ and $\mathbb{P}(\min\{\tau_A, \tau_B\} < \infty) = 1$.

The naive approach to estimating ℓ via Monte Carlo is to simulate a large number, N, of realizations of $\{X_n\}_{n\in\mathbb{N}_0}$ and count the number of times that a realization hits A before B. This leads to the *crude Monte Carlo* (CMC) estimator

$$\widehat{\ell}_{\mathsf{CMC}} = \frac{1}{N}\sum_{i=1}^{N} I_i,$$

where the $\{I_i\}_{i=1}^N$ are independent replicates of the random variable $\mathbb{I}(\tau_A < \tau_B)$. The efficiency of such an estimator is typically measured by its relative error, which is the standard deviation of the estimator divided by the quantity being estimated.

The idea of importance sampling is to produce a more efficient estimator by sampling the object of interest under a different probability measure, which makes the event of interest more likely. In the case of a Markov chain, we do this as follows. We define a new transition matrix, Q, by considering products of the form

$$Q(x, y) = r(x, y)P(x, y) \quad \text{for all } x, y \in \mathcal{X},$$

where $\sum_{y\in\mathcal{X}} Q(x, y) = 1$ for all $x \in \mathcal{X}$ and $r(x, y) > 0$ for all $x, y \in \mathcal{X}$ such that $P(x, y) > 0$ and $y \notin B$. The importance sampling estimator is then defined as follows. Independent replicates, $\{Y_n^{(1)}\}_{n\in\mathbb{N}_0}, \ldots, \{Y_n^{(N)}\}_{n\in\mathbb{N}_0}$, of a Markov chain starting at x_0 with transition matrix Q are simulated. Then,

$$\widehat{\ell}_{\mathsf{IS}} = \sum_{i=1}^{N} L\left(\{Y_n^{(i)}\}_{n\in\mathbb{N}_0}\right)\mathbb{I}(\tau_A^{(i)} < \tau_B^{(i)}) \tag{1}$$

is calculated, where $\tau_A^{(i)} = \inf\{n \in \mathbb{N}_0 : Y_n^{(i)} \in A\}$, $\tau_B^{(i)} = \inf\{n \in \mathbb{N}_0 : Y_n^{(i)} \in B\}$, $\tau^{(i)} = \min\{\tau_A^{(i)}, \tau_B^{(i)}\}$ and

$$L\left(\{Y_n^{(i)}\}_{n\in\mathbb{N}_0}\right) = \prod_{n=0}^{\tau^{(i)}} \frac{P(Y_n^{(i)}, Y_{n+1}^{(i)})}{Q(Y_n^{(i)}, Y_{n+1}^{(i)})} = \prod_{n=0}^{\tau^{(i)}} \frac{1}{r(Y_n^{(i)}, Y_{n+1}^{(i)})}.$$

The path functional $L(\cdot)$, called the likelihood ratio, weights the samples to ensure an unbiased estimator; see, e.g., L'Ecuyer et al. (2009). If the new transition matrix, Q, is chosen well, the relative error of the importance sampling estimator will be substantially smaller than that of the CMC estimator.

In principle, it is possible to construct a zero-variance change of measure in this setting using Doob's h-transform; see Asmussen and Glynn (2007). The resulting Markov chain has the same dynamics as the original chain conditioned on the event $\{\tau_A < \tau_B\}$. Although this is not practically useful (because it requires knowledge of the very probabilities that we are trying to estimate) it does give insight into what a good change of measure should look like. For more information on applying importance sampling to Markov chains, see Kroese et al. (2011) and Asmussen et al. (2013).

4 APPLYING IMPORTANCE SAMPLING TO AN INTERACTING PARTICLE SYSTEM

Applying importance sampling to interacting particle systems is challenging for a number of reasons. In particular, it is hard to find simple changes of measure that can make the event of interest more likely. For example, in the setting of Section 2, it may not be wise to 'push' the electron towards the set A, as this could also push the electron in the direction of holes. Even if there are clear ways to make the event of interest more likely, these may be atypical under the original measure. Furthermore, the state spaces of interacting particle systems are high dimensional and the change of measure clearly needs to be state dependent. This makes it difficult to use approaches that attempt to learn good changes of measure, such as the cross-entropy method (see Rubinstein and Kroese (2004)).

The heuristic that we adopt is to make only local changes. We evolve the system as normal except when the electron gets too close to a hole. When this occurs, we change the transition probabilities so that a collision between the electron and the hole is less likely to occur. This approach has a number of advantages. Firstly, we only need to update the likelihood ratio and calculate new probabilities when the electron gets too close to a hole. Secondly, because we do not change the transition probabilities very often, we do not sample under a measure that is dramatically different to the original measure. Thirdly (and, arguably, most importantly), when

we do apply importance sampling in a step of the algorithm, the likelihood ratio is multiplied by a factor that is strictly less than 1. This means that there is no risk of variance blow-up and, ignoring issues surrounding the additional computation time, we do not perform worse than CMC.

There are a number of different ways to reduce the probability that an electron collides with a nearby hole. Due to space issues, we describe only one approach here. In this approach, the electron and holes move less frequently to vertices occupied by the opposite type of particle than they would under the original probability measure. The frequency with which they move is controlled by a parameter $\rho \in (0, 1]$. We set

$$q^E(v \,|\, e, h^{(1)}, \ldots, h^{(K)}) = \frac{w^E(v \,|\, e, h^{(1)}, \ldots, h^{(K)})}{C(e, h^{(1)}, \ldots, h^{(K)})} \cdot p^E(v \,|\, e, h^{(1)}, \ldots, h^{(K)}),$$

where

$$w^E(v \,|\, e, h^{(1)}, \ldots, h^{(K)}) = \begin{cases} \rho & \text{if } v = h^{(k)} \text{ for some } k \in \{1, \ldots, K\}, \\ 1 & \text{otherwise.} \end{cases}$$

Similarly, we set

$$q^H(k, v \,|\, e, h^{(1)}, \ldots, h^{(K)}) = \frac{w^H(v \,|\, e, h^{(1)}, \ldots, h^{(K)})}{C(e, h^{(1)}, \ldots, h^{(K)})} \cdot p^H(k, v \,|\, e, h^{(1)}, \ldots, h^{(K)}),$$

where

$$w^H(v \,|\, e, h^{(1)}, \ldots, h^{(K)}) = \begin{cases} \rho & \text{if } v = e \\ 1 & \text{otherwise.} \end{cases}$$

The probabilities are normalized so that they sum to one. That is, we take

$$\begin{aligned} C(e, h^{(1)}, \ldots, h^{(K)}) = &\sum_{v \in V} w^E(v \,|\, e, h^{(1)}, \ldots, h^{(K)}) \cdot p^E(v \,|\, e, h^{(1)}, \ldots, h^{(K)}) \\ &+ \sum_{k=1}^{K} \sum_{v \in V} w^H(v \,|\, e, h^{(1)}, \ldots, h^{(K)}) \cdot p^H(k, v \,|\, e, h^{(1)}, \ldots, h^{(K)}). \end{aligned}$$

Let $\{\check{S}_n\}_{n \in \mathbb{N}_0} = \{(\check{E}_n, \check{H}_n^{(1)}, \ldots, \check{H}_n^{(K)})\}_{n \in \mathbb{N}_0}$ be the Markov chain defined by these dynamics. Note that, when no hole is adjacent to the electron, the dynamics of $\{\check{S}_n\}_{n \in \mathbb{N}_0}$ are exactly the same as those of the original chain. The functions $\frac{w^H(v \,|\, e, h^{(1)}, \ldots, h^{(K)})}{C(e, h^{(1)}, \ldots, h^{(K)})}$ and $\frac{w^E(v \,|\, e, h^{(1)}, \ldots, h^{(K)})}{C(e, h^{(1)}, \ldots, h^{(K)})}$ play the role of $r(x, y)$ in the importance sampling estimator given in (1) and are used to construct the likelihood ratio. A replicate of $L(\{\check{S}_n\}_{n \in \mathbb{N}_0}) \mathbb{I}(\tau_{\mathsf{elec}} < \tau_{\mathsf{recomb}})$ can be simulated as follows.

Algorithm 4.1 (Importance sampling by changing the jump frequencies)**.**

1. Set $\check{S}_0 = (e_0, h_0^{(1)}, \ldots, h_0^{(K)})$. Set $n = 0$. Set $L_0 = 1$.
2. Draw $U \sim \mathcal{U}(0, 1)$.
3. If $U \leq \sum_{v \in V} q^E(v \,|\, \check{S}_n)$, move the electron. That is, draw $\check{E}_{n+1} \sim \mu^E(\cdot \,|\, S_n)$, where

$$\mu^E(v \,|\, \check{S}_n) = \frac{q^E(v \,|\, \check{S}_n)}{\sum_{v' \in V} q^E(v \,|\, \check{S}_n)} \quad \text{for all } v \in V.$$

   Set $H_{n+1}^{(k)} = H_n^{(k)}$ for all $k \in \{1, \ldots, K\}$. Set

$$L_{n+1} = \frac{C(\check{S}_n)}{w^E(\check{E}_{n+1} \,|\, \check{S}_n)} \cdot L_n.$$

4. If $U > \sum_{v \in V} q^E(v \,|\, \check{S}_n)$, move a hole. That is, draw $k^* \sim \mu^K(\cdot \,|\, \check{S}_n)$, where

$$\mu^K(k \,|\, \check{S}_n) = \frac{\sum_{v \in V} q^H(k, v \,|\, \check{S}_n)}{\sum_{k'=1}^{K} \sum_{v \in V} q^H(k', v \,|\, \check{S}_n)} \quad \text{for all } k \in \{1, \ldots, K\}.$$

Then, draw the new state, $\check{H}_{n+1}^{(k^*)} \sim \mu^H(\cdot \mid k^*, S_n)$, where

$$\mu^H(v \mid k^*, S_n) = \frac{q^H(k^*, v \mid \check{S}_n)}{\sum_{v' \in V} q^H(k^*, v' \mid \check{S}_n)} \quad \text{for all } v \in V.$$

Set $\check{E}_{n+1} = \check{E}_n$ and $\check{H}_{n+1}^{(k)} = \check{H}_n^{(k)}$ for all $k \in \{1, \ldots, K\} \setminus \{k^*\}$. Set

$$L_{n+1} = \frac{C(\check{S}_n)}{w^H(\check{H}_{n+1}^{(k^*)} \mid \check{S}_n)} \cdot L_n.$$

5. If $\check{E}_n = \check{H}_n^{(k)}$ for some $k \in \{1, \ldots, K\}$, return 0. If $\check{E}_n \in A$, return L_n. Otherwise, set $n = n + 1$ and repeat from step 2.

Note that, in practice, $C(\check{S}_n)$ is straightforward to compute and $\mu^E(\cdot \mid \check{S}_n)$, $\mu^K(\cdot \mid \check{S}_n)$ and $\mu^H(\cdot \mid k^*, S_n)$ can be easily sampled from via acceptance-rejection.

5 EXPERIMENTS

For reasons of space, we only consider one specific example of the problem here. The basic setup, illustrated in Figure 2, is the following. We consider an $m \times m$ square lattice. The lattice has periodic boundary conditions above and below and reflective boundary conditions to the left and right. The initial configuration of the system consists of $m \times l$ holes, which occupy the rightmost l columns. The electron starts in the leftmost column. The electrode is the rightmost column. With probability α, the electron moves right. Otherwise, it selects one of the other directions uniformly at random. With probability β, a hole attempts to move left. Otherwise, it selects one o f the other directions uniformly at random. The parameter l allows us to control the rarity of the event $\{\tau_{\mathsf{elec}} < \tau_{\mathsf{recomb}}\}$.

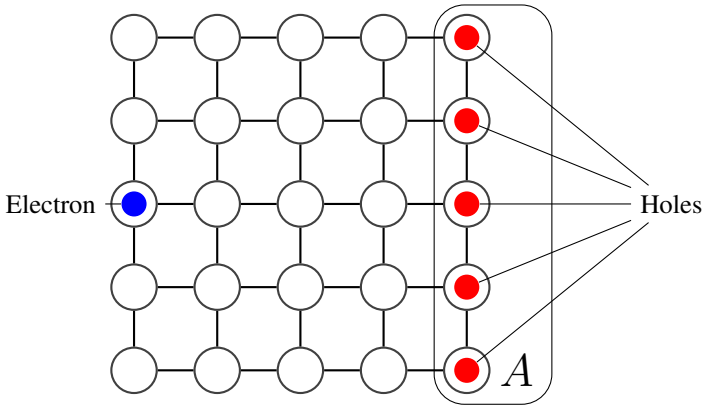

Figure 2. Illustration of the setup used in the numerical experiments, with $m = 5$ and $l = 1$.

The results of the experiments, with $m = 10$, $\alpha = .7$ and $\beta = .7$, are shown in Table 1. For each l, the best case relative error of the importance sampling estimator is roughly $1/2$ of the relative error of the corresponding CMC estimator. This means that the importance sampling estimator can obtain a similar level of accuracy to the CMC estimator using $1/4$ as many samples. Although the reduction in relative error is not a large improvement in the classical rare event simulation context, it is a significant improvement in our context. This is because simulations of realistic physical systems often need to be run over periods of several weeks (and, thus, reducing the sample size by a factor of 2–4 is significant). We observe that, although the importance sampling estimator require more computational effort per sample, this overhead can be mitigated through careful implementation.

Table 1. Estimates ($\widehat{\ell}$) and estimated relative errors ($\widehat{\mathsf{RE}}$) for the crude Monte Carlo (CMC) and importance sampling (IS) estimators with sample size $N = 10^7$.

	CMC	IS($\rho = 0.01$)	IS ($\rho = 0.05$)	IS ($\rho = 0.10$)	IS ($\rho = 0.50$),
$l = 1$					
$\widehat{\ell}$	1.730×10^{-1}	1.729×10^{-1}	1.729×10^{-1}	1.731×10^{-1}	1.731×10^{-1}
$\widehat{\mathsf{RE}}$	6.15×10^{-4}	3.78×10^{-4}	3.87×10^{-4}	3.96×10^{-4}	4.81×10^{-4}
$l = 2$					
$\widehat{\ell}$	2.114×10^{-2}	2.130×10^{-2}	2.108×10^{-2}	2.115×10^{-2}	2.113×10^{-2}
$\widehat{\mathsf{RE}}$	2.13×10^{-3}	1.21×10^{-3}	1.23×10^{-3}	1.26×10^{-3}	1.55×10^{-3}
$l = 3$					
$\widehat{\ell}$	1.784×10^{-3}	1.787×10^{-3}	1.783×10^{-3}	1.785×10^{-3}	1.781×10^{-3}
$\widehat{\mathsf{RE}}$	7.47×10^{-3}	3.33×10^{-3}	3.42×10^{-3}	3.53×10^{-3}	4.70×10^{-3}

6 EXTENSIONS

The methodology presented above represents only the first steps on the road to finding effective importance sampling strategies for interacting particle systems. New techniques need to be developed and the existing strategies need to be tested in a number of different settings to identify where they work well and where they fail. Currently, the authors are working on changes of measure that consider what may occur several steps ahead. This allows for strategies where holes organize to let the electron pass and where the electron is able to start moving away from a hole that is not immediately adjacent to it. These 'locally optimal' changes of measure can be constructed by solving discrete Dirichlet problems, though the dimension of these problems can be prohibitively high.

REFERENCES

Asmussen, S., P. Dupuis, R. Rubinstein, and H. Wang (2013). Importance sampling for rare events. In M. Fu and S. Gass (Eds.), *Encyclopedia of Operations Research and Management Sciences* (Third ed.). Springer.

Asmussen, S. and P. W. Glynn (2007). *Stochastic Simulation*. New York: Springer.

Brereton, T., O. Stenzel, B. Baumeier, D. Andreinko, V. Schmidt, and D. P. Kroese. (2014). Efficient simulation of Markov chains using segmentation. *Methodology and Computing in Applied Probability 16*, 465–484.

Grimmett, G. (2010). *Probability on Graphs*, Volume 1 of *Institute of Mathematical Statistics Textbooks*. Cambridge University Press, Cambridge.

Kroese, D. P., T. Taimre, and Z. I. Botev (2011). *Handbook of Monte Carlo Methods*. New York: John Wiley & Sons.

L'Ecuyer, P., M. Mandjes, and B. Tuffin (2009). Importance sampling in rare event simulation. In G. Rubino and B. Tuffin (Eds.), *Rare Event Simulation using Monte Carlo Methods*, pp. 17–38. New York: John Wiley & Sons.

Liggett, T. M. (2005). *Interacting Particle Systems*. Classics in Mathematics. Springer-Verlag, Berlin. Reprint of the 1985 original.

Rubinstein, R. Y. and D. P. Kroese (2004). *The Cross-Entropy Method: A Unified Approach to Combinatorial Optimization Monte-Carlo Simulation, and Machine Learning*. New York: Springer.

Stenzel, O., C. Hirsch, V. Schmidt, T. Brereton, D. P. Kroese, B. Baumeier, and D. Andrienko (2014). A general framework for consistent estimation of charge transport properties via random walks in random environments. *Multiscale Modeling and Simulation 12*, 1108–1134.