

An Analytical Study of Monte Carlo Source Term Estimators in Plasma Edge Simulations of Fusion Reactors

B. Mortier*, M. Baelmans[†], G. Samaey*

*Department of Computer Science, KU Leuven, Celestijnenlaan 200A, 3001 Leuven, Belgium

[†]Department of Mechanical Engineering, KU Leuven, Celestijnenlaan 300, 3001 Leuven, Belgium
bert.mortier@cs.kuleuven.be

Abstract - The statistical error and computational efficiency of Monte Carlo reaction rate estimation depends crucially on which Monte Carlo procedures are used. We present a systematic comparison of the variance and computational cost of the most common estimation procedures and several others, and we show a non-trivial dependence of the choice of the optimal estimator on the problem parameters. The comparison is based on an invariant imbedding procedure, in which systems of ordinary differential equations (ODEs) are derived that quantify the statistical error and computational cost of each estimator. We analyze in detail a scenario with forward-backward scattering in a one-dimensional slab. For this test case, we perform a parametric study of the expected statistical error and computational cost by numerically solving the ODEs.

I. INTRODUCTION

Plasma edge simulations are important in the analysis and design of the heat and particle exhaust in nuclear fusion divertor analysis and design [1]. These simulations have to take into account the behavior of plasma and neutral particles. The plasma is usually well described by Navier-Stokes-like equations, the Braginskii equations, that can be simulated with a finite volume (FV) code. The neutral particle distribution requires modeling in a position-velocity phase space with kinetic equations. The additional dimensions give a Monte Carlo (MC) method the upper hand over FV.

Coupling both subsimulations is a challenge due to the stochastic nature of the MC method. Plasma and neutral particles interact through (at least) absorption and scattering collisions. These interactions cause the exchange of mass, momentum and energy between the plasma and the neutral particles. Hence, the stochasticity of the MC neutral simulation is transferred to the FV plasma simulation. The resulting statistical noise can hamper the convergence of the simulation. To reduce the impact of the noise, the statistical error of the MC simulation has to be reduced. This can be done by either increasing the number of particles, and consequently the computational cost, or by reducing the variance on the simulation of a single neutral particle.

A fundamental way of reducing the variance on the simulation of a single MC particle is selecting an estimator, which attains a lower statistical error on the quantities of interest for the same computational cost [2]. Which estimation procedure is most efficient depends non-trivially on the problem under study. To make a well-founded choice, a sound understanding of the estimators' behavior for different problems is required, as well as quantitative guesses of potential gains in changing the estimating procedure. We assess the effect of different choices in the MC estimators by conducting a parametric study for a coherent and relevant set of estimators. To this end, we construct ODEs that provide the variance of the different estimators in a one-dimensional slab setting. These are then numerically simulated for different values of the scattering and absorption rates, as well as the anisotropy, i.e., direction-preference of the post-collision velocity.

In [3], MacMillan derives ODEs for the statistical error of a small set of reaction rate estimation procedures in a forward scattering scenario. Indira [4] similarly studied the statistical error of several for leakage estimation procedures, but included forward-backward scattering.

Both the setting in [3] and [4] allowed for significant simplifications in the ODEs for the statistical error, which have a comprehensive analytical solution. In contrast, the model problem with forward-backward scattering we consider for reaction rate estimation is much more relevant to realistic plasma edge simulations. This results in much larger systems of ODEs, which need to be solved numerically.

Other analytical studies with a similar aim of analytically determining the variance of the most-used estimation procedures, are most notable by Lux. He determined approximate formulas for the variance of different estimators [5], and sufficient conditions for one estimator to outperform another [6]. We also refer to [7–11] which use analytical calculations of the variance for variance reduction through importance sampling or similar techniques.

In section II we describe the model and the problem under study. Subsequently, the estimation procedures are shortly discussed: in section III we introduce the most-used estimators in reaction rate estimations and in section IV we explain how the simulation itself can be adapted to reduce variance. In section V we illustrate how the systems of ODEs are derived via an invariant imbedding procedure. Section VI contains a comparison of the performance of the different estimation procedures throughout the parameter space. The potential application of these results are discussed in section VII. Finally, a short summary and future prospects are given in section VIII.

II. MODEL SETUP AND ANALOG MC SIMULATION

In a practical plasma edge simulation, the simulated neutral particles move in a 3D/3D (position \mathbf{r} /velocity \mathbf{v}) space. The particles originate at the target by recombination of the plasma, indicated by a source term $S(\mathbf{r}, \mathbf{v})$, which is only nonzero at the target. As they move, neutral particles undergo two types of collisions: charge-exchange collisions, during which a neutral particle collides with a plasma ion and

they exchange charge, and ionization, during which a neutral particle gets absorbed by the plasma. Since plasma particles are not modelled explicitly, charge-exchange is simulated by a scattering event, in which a neutral particle is removed from the simulation and replaced by a new neutral particle with velocity drawn from the distribution $f_{\nu \text{ postcol},r}(\nu)$, which depends on the plasma. The rates at which each of these collision types occur are given by the cross-sections $\Sigma_s(\mathbf{r}, \nu)$ and $\Sigma_a(\mathbf{r}, \nu)$. The resulting stationary kinetic equation for the neutral particle density distribution function $\phi_n(\mathbf{r}, \nu)$ is

$$\underbrace{\nu \cdot \nabla_{\mathbf{r}} \phi_n(\mathbf{r}, \nu)}_{\text{movement}} = \underbrace{S(\mathbf{r}, \nu)}_{\text{source from the plasma}} - \underbrace{\Sigma_s(\mathbf{r}, \nu)|\nu| \phi_n(\mathbf{r}, \nu)}_{\text{sink due to scattering}} - \underbrace{\Sigma_a(\mathbf{r}, \nu)|\nu| \phi_n(\mathbf{r}, \nu)}_{\text{sink due to absorption}} + \underbrace{f_{\nu \text{ postcol},r}(\nu) \int_{-\infty}^{+\infty} \Sigma_s(\mathbf{r}, \nu') |\nu'| \phi_n(\mathbf{r}, \nu') d\nu'}_{\text{source due to scattering}} .$$

Studying the variance of source term estimators in this general 3D/3D setting quickly becomes intractable due to the large dimension of the parameter space that needs to be considered. Here, we therefore restrict to a one-dimensional domain of length L . Additionally, the cross-sections for absorption (ionization) Σ_a and scattering (charge-exchange) Σ_s are taken to be constant in this domain. Finally, we consider the velocity to always be ± 1 . All particles enter the one-dimensional domain from the left. Then, the only additional parameter that remains is the probability P_r that the particle has a post-collision velocity of $+1$. In that sense, the deviation of P_r from 0.5 represents the anisotropy of scattering collisions.

To simulate a particle, we give it an initial position at the left boundary ($x = 0$) and an initial velocity 1. A distance to the next collision is then sampled from the exponential distribution $\rho \propto \mathcal{E}(\Sigma_t)$ with $\Sigma_t = \Sigma_a + \Sigma_s$. If, after this distance, the particle would be outside of the domain, no collision occurs and the particle is taken out of the simulation. If not, there is a probability of Σ_a/Σ_t the particle gets absorbed and disappears from the simulation. With probability Σ_s/Σ_t the particle undergoes a scattering event and a new velocity is sampled from $\{-1, 1\}$ with corresponding probabilities $\{1 - P_r, P_r\}$. New collision positions are sampled until the particle disappears from the simulation. Such a simulation tries to simulate the particle as if it is a physical particle. These simulations are called *analog*, and we denote them by \mathbf{a} .

This very simple model does maintain a strong connection to actual neutral particle simulations. It can be viewed as the simulation of one passage of a particle through a grid cell of the piecewise constant discretization of the plasma variables, albeit with reduced dimensionality.

III. ESTIMATORS

In our study, we will focus on Monte Carlo simulations of which the goal is to estimate the expected number of collisions of a certain type a neutral particle undergoes when entering the 1D slab (or the grid cell) from the left. If, for instance, the expected exchange of mass with the plasma were to be the

goal of the simulation, this number should be multiplied with the average amount of exchanged mass per collision.

The most commonly used reaction rate estimators are collision estimators, track-length estimators and next-event estimators (also called expectation estimators) [12]. These estimators will be introduced briefly, but first we discuss analog estimators, which are more fundamental.

Analog Estimator. The basic way to estimate a source term from a Monte Carlo simulation is by simply counting what happens, as it happens in the physically based simulation. To estimate the expected number of absorption collisions, for instance, one scores 1 every time an absorption occurs. This analog estimator for absorption will be denoted by $\mathbf{a_abs}$. When counting scattering events, we denote the analog estimator by $\mathbf{a_sc}$.

Collision Estimator. The collision estimator, as well as the subsequent estimators, cease counting what physically happens, while remaining unbiased. The idea of the collision estimator is to count the probability the collision turns out to be of the type to be estimated, at every collision. To estimate the amount of absorption events, one scores Σ_a/Σ_t at every collision, irrespective of whether it is an absorption or a scattering collision. If one estimates the expected number of scattering collisions, Σ_s/Σ_t is scored at every collision. Similar as for the analog estimators, we denote the two kinds of collision estimators by $\mathbf{c_abs}$, respectively $\mathbf{c_sc}$. For the collision estimator, nor for the next-event and track-length estimators, which are explained in the subsequent paragraphs, it is necessary to distinguish between the type of event of interest. For these three, this would only amount to a scaling of the scores, which has no impact on our performance measures that look at the relative statistical error. In the remainder of this text we will denote by Σ the cross-section of the collision type of interest.

Next-Event Estimator. As the name indicates, the estimator looks as far as the next collision. It counts the probability that a next collision of the type of interest occurs within the domain. Since the collision events are exponentially distributed, the probability for a new collision before the end of the domain is

$$\int_0^D \Sigma_t e^{-\Sigma_t \ell} d\ell = 1 - e^{-\Sigma_t D} ,$$

with D the distance from the current collision point to the end of the domain, given that the particles continues in its post-collision velocity. The probability of that next collision being of the type of interest is Σ/Σ_t . So at every collision, this estimator counts

$$\frac{\Sigma}{\Sigma_t} \int_0^D \Sigma_t e^{-\Sigma_t \ell} d\ell = \frac{\Sigma}{\Sigma_t} (1 - e^{-\Sigma_t D}) . \quad (1)$$

We denote the next-event estimators by $\mathbf{ne_abs}$ and $\mathbf{ne_sc}$.

Track-Length Estimator. The track-length estimator can be derived by partial integration of Eq. (III):

$$\begin{aligned} \frac{\Sigma}{\Sigma_t} \int_0^D \Sigma_t e^{-\Sigma_t \ell} d\ell &= \Sigma \ell e^{-\Sigma_t \ell} \Big|_0^D + \int_0^D \Sigma \Sigma_t \ell e^{-\Sigma_t \ell} d\ell \\ &= \Sigma \left(D e^{-\Sigma_t D} + \int_0^D \Sigma_t \ell e^{-\Sigma_t \ell} d\ell \right) . \end{aligned}$$

The two terms between the brackets equal the expected distance the particle will travel in the domain without colliding. There is a probability of $e^{-\Sigma_t D}$ that it will not collide, with D the distance travelled in that case. The second term between brackets is the integral over the product of a distance ℓ times the probability of colliding for the first time at that distance ℓ . This means another unbiased estimator for the amount of collisions of the type of interest scores Σ times the distance travelled in the domain. This is the track-length estimator and the two versions denoted by `tl_abs`, respectively `tl_sc`.

IV. NON-ANALOG MC SIMULATION

In contrast to the estimator type, which only changes what is scored during the simulation, *survival biasing* (also called implicit capture) impacts the simulation itself. The particles no longer behave as physical particles; hence the name non-analog. The aim is to let particles penetrate deeper into the domain, so less frequented areas are still sampled sufficiently [12]. In concreto, survival biasing entails that absorption is not executed by letting the particle disappear, but by reducing its weight with the probability of absorption. This way, particles are much more likely to traverse regions with a high absorption probability, a ‘survival bias’ that is compensated by a decrease of the particle weight. We investigate the effects of two options to assign the absorption events: to the collision events or along the path, these form the two limiting cases as described by [13, 14]. Except for the analog estimators, the described estimators can be used in an analog simulation and both of these non-analog simulations. Several source term estimation procedures have been conceived based on non-analog MC simulation. We briefly review the two most common ones, for which we introduce a name that allows treating them in a unified way with the estimators of section III. To allow for a systematic comparison and discussion we introduce a terminology that connects the estimator types to the simulation types.

Non-Analog Collision Type Simulation. In the analog situation, there is a probability of Σ_a/Σ_t to be absorbed at a collision (which means the particle is removed from the simulation) and a probability of Σ_s/Σ_t to be scattered. One can always simulate a scattering event at a collision, while simultaneously multiplying the weight by Σ_s/Σ_t to account for the probability that the event was an absorption. Due to its focus on the collision events this type of simulation will be called the *non-analog collision type simulation* and is denoted by `nac`.

Non-Analog Track-Length Type Simulation. The second option is to have the absorption events be carried out along the path. When the particle undergoes a scattering event after a certain distance d , one can change its weight with the expected fraction of particles that were to be absorbed before the scattering location was reached. This means only the collision density of scattering events is sampled, so this should be taken into account while scoring. This *non-analog track-length type simulation* will be denoted by `natl`.

There are several ways to understand these methods. The simplest interpretation is probably by considering each simulated particle as to represent an infinite number of ‘real’ parti-

cles that move coherently and collide simultaneously. During this collision (for `nac`) or along their trajectory (for `natl`), a fraction of these particles get absorbed, whereas the remaining fraction continues to move coherently with the newly selected velocity.

V. INVARIANT IMBEDDING PROCEDURE

The previously discussed estimators (`c_abs`, `c_sc`, `tl_abs`, `tl_sc`, `ne_abs` and `ne_sc`), combined with the different simulation types (`a`, `nac`, `natl`), give 18 estimation procedures, which we supplement with the basic estimators `a_abs` and `a_sc` for the analog simulation. We study their performance under different values for the parameters of the simplified problem discussed in section II (Σ_a , Σ_s , P_r). Their performance is measured in both the relative statistical error and the computational cost, given a certain statistical error.

To make a quantitative comparison of these performance measures in the entire three-dimensional parameter space, MC experiments are not a feasible option. Estimating variance sufficiently accurately would be too computationally expensive for a three-dimensional set of parameter values. Instead, an invariant imbedding procedure is used to derive ODEs that can be used to calculate the statistical error on the computed number of collisions, as well as the computational cost associated with a fixed statistical error.

The invariant imbedding procedure [15, 16] for the variance consists of writing the first and second moment of the score in a slab of length $x + \Delta x$ with fixed parameters (Σ_s , Σ_a , P_r) in terms of the scores for a slab of length x . By taking the limit of $\Delta x \rightarrow 0$, an ODE is formed with the domain length as an integration variable. For the other quantities that pop up, a similar procedure can be followed until the set of ODEs is closed.

As an example, we will elaborate an invariant imbedding procedure for the second moment of a non-analog collision type track-length estimation procedure (`nac_tl_abs` or `nac_tl_sc`) under the condition that the particle leaves the domain, the same way as it entered. For the detailed derivations and the other scenarios, we refer to [17].

Invariant imbedding example

To recapitulate, a non-analog collision type simulation (`nac`) executes every collision as a scattering event, keeping it unbiased by reweighing the particle by a factor Σ_s/Σ_t at every collision. A track-length estimator (`tl_abs` or `tl_sc`) counts $\Sigma \ell$, with ℓ the distance travelled and Σ the cross-section of the scored event, which can be either Σ_a or Σ_s .

Let T denote an arbitrary `nac_tl_abs` or `nac_tl_sc` estimation. We will add a subscript to indicate if a particle enters the domain from the left (l) or from the right (r), and similarly for exiting the domain. For instance, T_{ll} denotes an arbitrary `nac_tl_abs` or `nac_tl_sc` estimation with the condition that the particle leaves and enters the domain from the left. Its expected value is denoted by t_{ll} and its second moment $\mathbb{E}[T_{ll}^2]$ by τ_{ll} . We can group all the possible particle paths that enter and leave from the left side in a domain of length $x + \Delta x$ based on the behavior inside the leftmost part

of length Δx . We neglect paths that collide more than once in the Δx part of the domain, since these would occur with a probability of order $O(\Delta x^2)$, $\Delta x \rightarrow 0$. There are five options left, which are shown in Fig. 1. The behavior of the particle in the rightmost part of the domain with length x is arbitrary, as long as the particle enters that part of the domain from the left, and leaves it to the left again. So the amount and location of collisions inside that part of the domain is irrelevant, besides that condition. The symbol \bullet in the part of length Δx identifies that a collision took place there. After the collision, the particle goes right with a probability of P_r , and goes left with a probability of $1 - P_r$.

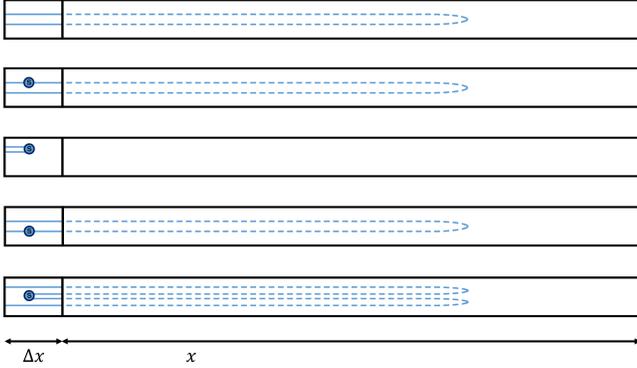


Fig. 1: The possible paths in a domain of length $x + \Delta x$ that start and end on the left side in a non-analogous simulation and have a probability of order at most 1 in Δx .

For each of these five situations, we can write its probability of occurring and the expected second moment given that it occurs in terms of quantities for a domain of length x . For instance, the first situation occurs with a probability of $P_{ll}(x)(1 - e^{-\Sigma_t \Delta x})^2$, with $P_{ll}(x)$ the probability of a particle that enters from the left in a domain of length x to also leave at the left side and $(1 - e^{-\Sigma_t \Delta x})$ is the probability of a particle not reacting during a movement through a slab of length Δx . Only retaining first order terms, this probability can be rewritten as $P_{ll}(x)(1 - 2\Sigma_t \Delta x)$. The expected second moment of such a path equals

$$\mathbb{E}[(T_{ll}(x) + \Sigma \Delta x + W_{ll} \Sigma \Delta x)^2] = \tau_{ll}(x) + 2\Sigma \Delta x t_{ll}(x) + 2\Sigma \Delta x \theta_{ll} + O(\Delta x^2), \Delta x \rightarrow 0.$$

In this equation W_{ll} stands for the weight reduction of a particle after a passage through a domain of length x to which it entered and returned from the left and $\theta_{ll} = \mathbb{E}[W_{ll} T_{ll}]$. The product of both the probability and its expected second moment constitutes the contribution of the paths with such a behavior to $\tau_{ll}(x)$. For the first situation of Fig. 1 this equals $P_{ll}(x)(\tau_{ll}(x) - 2\Sigma_t \Delta x \tau_{ll}(x) + 2\Sigma \Delta x t_{ll}(x) + 2\Sigma \Delta x \theta_{ll}(x))$, up to order 1 in Δx .

For the other four cases of Fig. 1 one can work similarly, taking into account that at collisions the particle is reweighed by a factor Σ_s / Σ_t . In this process one final new variable arises, namely $\omega_{ll} = \mathbb{E}[W_{ll}^2]$. The result for each of the five cases separately amounts to

$$1. P_{ll}(x)(\tau_{ll}(x) - 2\Sigma_t \Delta x \tau_{ll}(x) + 2\Sigma \Delta x t_{ll}(x) + 2\Sigma \Delta x \theta_{ll}(x))$$

2. $P_r \frac{\Sigma_s}{\Sigma_t} \Delta x P_{ll}(x) \tau_{ll}(x)$
3. 0
4. $P_{ll}(x) \Sigma_t \Delta x (1 - P_r) \tau_{ll}(x)$
5. $P_{ll}^2(x) \Delta x P_r (\Sigma_t \tau_{ll}(x) + 2\Sigma_s t_{ll}(x) \theta_{ll}(x) + \frac{\Sigma_s}{\Sigma_t} \omega_{ll}(x) \tau_{ll}(x))$.

Adding these terms together gives an expression for $P_{ll}(x + \Delta x) \tau_{ll}(x + \Delta x)$, which, by taking the limit $\Delta x \rightarrow 0$, can be transformed into an ordinary differential equation for $P_{ll} \tau_{ll}$. For each of the other variables in this equations (P_{ll} , $P_{ll} t_{ll}$, $P_{ll} \omega_{ll}$, $P_{ll} \theta_{ll}$), this process can be repeated. Only one extra variable pops up in this process, namely $P_{ll} w_{ll}$, with $w_{ll} = \mathbb{E}[W_{ll}]$. Again, this process can be followed and finally a closed set of differential equations can be found, for which the initial value is a zero-vector, because the probability of entering and leaving from the left in a slab of length zero is zero and P_{ll} features in each of the variables. The full details will be included in [17].

One last remark to be made is regarding the integration variable. We explained the method with the domain length as an integration variable, for fixed Σ_s , Σ_a and P_r . This would mean we need to integrate one set of ODEs up to L for each point in the parameter space we want to evaluate. Instead, we use the fact that for a fixed Σ_s / Σ_t , P_r and $\Sigma_t L$, a change in L does not change the simulation, nor the score. This means we can see $\Sigma_t L$ as the integration variable and we only need to integrate the ODEs as many times as we want different values of Σ_s / Σ_t and P_r . Each point in the integration of the ODEs from 0 to a maximal value of $\Sigma_t L$ serves then as a value of $\Sigma_t L$.

VI. RESULTS

For each of the estimation procedures, a system of ODEs has been derived and these are numerically simulated to evaluate both the statistical error and the computational cost. For the collision estimators (c_abs, c_sc), track-length estimators (t1_abs, t1_sc) and next-event estimators (ne_abs, ne_sc), both measures of performance are independent of the specific type of collision that is scored. The reason has been put forward in section III: changing the type of collision to tally only amounts to a scaling of the score. Hence, we can look at the _abs and _sc variants jointly. So in this section we denote the collision estimators by c, the track-length estimators by t1 and the next-event estimators by ne.

In part 1 of this section, we show and discuss the performance of the estimation procedures in terms of their relative statistical error. We focus on the dependence on the parameters of the problem and discuss the effects at play. Part 2 does the same for computational cost, which is a slightly more complex performance measure. Finally, in part 3, we make explicit which estimator forms the best choice in each part of the parameter space and show the potential loss of sticking to a single estimator only. We will discuss how these results generalise to higher dimensions in the discussion section.

1. Relative standard deviation

We first compare the different estimation procedures based on the relative standard deviation. If the number of particles entering the domain is given, the relative standard deviation is a measure for the relative statistical error.

A. Trivial scattering, $P_r = 1$

If $P_r = 1$, the scattering is trivial: the particles never change velocity since the post-scattering velocity is always 1, which is also the velocity at which the particles enter the simulation. This will enable a clear explanation of part of the mechanisms at hand. Still, even in this simple situation, it is nontrivial which of the mechanisms dominates for which parameter set, and what the quantitative results are.

This simple situation often has easily obtainable analytic solutions for the ODEs, as has been done by [3] for `a_a_abs`, `a_c`, `a_t1` and `a_ne`, as well as some combinations. Most of the results in this section are thus not new, but form a stepping stone towards the rest of our results which, to the best of our knowledge, are new.

Fig. 2 shows the results of this part. The axes are the non-dimensional scattering and absorption cross-sections of which combinations are taken up to $\Sigma_t L = 10$. For larger values of $\Sigma_t L$, the contour lines become invariant to L , so they become straight lines. This is due to the fact that with large collision probabilities, the probability of a particle reaching the end of the domain is very small, so their contribution becomes negligible. Since particles reaching the end constitute a negligible part, making the domain larger has no effect. In the discussion below, the different estimation procedures are indicated by pictograms that refer to the ordering of the combinations in the figure.

The analog absorption estimator (`a_a_abs`, ) is one of the most trivial estimators, both in how it works and in the resulting graph. The statistical error is independent of the scattering rate since the scattering events do not impact the particle, neither on its weight, nor on its trajectory. If $\Sigma_a L$ increases, the statistical error decreases, since it steadily becomes almost certain that the particle will be absorbed.

The analog track-length estimator (`a_t1`, ) is similar to `a_a_abs` in its independence to $\Sigma_s L$. Now the relative statistical error increases with rising $\Sigma_a L$. When there are hardly any absorptions, the path length, which scales with what is scored, is nearly always L , giving low variability. If $\Sigma_a L$ increases, absorptions occur more often, which randomly cut the path short, increasing the statistical error due to the increased stochasticity of the path length.

The stochasticity due to a path getting cut short prematurely when $\Sigma_a L$ increases is also visible in the analog scattering estimator (`a_a_sc`, ). When cut short early in the domain, a lot fewer contributing scattering collisions occur and vice versa for paths that are cut short much later than on average. Similarly increasing $\Sigma_s L$, lowers the relative statistical error due to the increased amount of scored events. These aspects also show themselves in the non-analog collision type track-length estimator (`nac_t1`, ). A higher absorption rate now does not mean the particle disappears from the simulation, but its weight gets lowered. With low survival probabilities

(Σ_s/Σ_t) this is nearly the same as being actually absorbed. With $\Sigma_s L = 0$ the results are the same as `a_t1`, but increasing $\Sigma_s L = 0$ now has a positive effect as it also had with the analog scattering estimator (). It namely increases the number of times the weight gets adapted, and as such lowers the variance on the path-averaged weight.

The analog collision estimator (`a_c`, ) combines traits of the analog scattering estimator () and the analog absorption estimator (). It scores the same at both scattering and absorption events, so increasing either $\Sigma_s L$ or $\Sigma_a L$, increases the amount of scored events. In some instances it decreases the variance and in others it increases. If scattering collisions are dominant, absorption has an adverse effect because it can make the amount of scattering collisions more variable. If absorption collisions are the dominant contributor to the score, increasing this further gives a near-certain absorption and as such a low variance. Similarly, if there is nearly certainly an absorption event going to happen, scattering events introduce variance on the number of events that was not yet there.

In the non-analog collision type collision estimator (`nac_c`, ) increasing the absorption rate never has an adverse effect. The only thing that is now stochastic, is the amount of collisions, which all have the same effect. The picture is not symmetric however, due to the fact that the weight is multiplied by the survival probability after every collision. The lower the survival probability, the less significant the later collisions are, and since there is more variability in the occurrence of these, an increase of $\Sigma_a L$ is more positive for the variance than an increase of $\Sigma_s L$.

In the non-analog track-length type collision estimator (`nat1_c`, ) the absorption events are no longer used as scoring events, so only the negative effects are retained. In some parts of the domain it does perform better than the analog collision estimator () because the randomness due to the sudden or very late path cut-off is gone.

The non-analog track-length type track-length estimator (`nat1_t1`, ) stands out here, since it has zero variance throughout the domain. It scores based on a certain total path length, and since the path length is always L in non-analog simulations with only trivial scattering, this estimator performs perfectly. A good comparison to be made here is with the non-analog track-length type next-event estimator (`nat1_ne`, ). As also stated by Lux [12], it is natural to — wrongly — expect the next-event estimator to outperform all others like the collision and the track-length estimators. Comparing ) and ) shows this is indeed not true. Albeit the next-event estimator can be written as scoring the expected contribution of a track-length estimator for the next collision. In this simple case it is easy to see why this happens. The total path is not stochastic due to the fact that $P_r = 1$, but the expected path length until the next collision is.

Something non-trivial that pops up in each of the next-event estimators () is the perfect score when $\Sigma_s L = 0$. An increase in $\Sigma_s L = 0$ first increases the relative standard deviation. Increasing $\Sigma_s L$ further then decreases the relative standard deviation again. The reason is that without scattering events a next-event estimator counts exactly once: at entry of the domain. If scattering events do occur, the estimator also scores at them with a variable score depending on where the

scattering event occurred exactly. The initially stochasticity-free scoring now has a variable aspect, hence the increase in statistical error. For high enough $\Sigma_s L$, the relative standard deviation decreases again due to the increase in scored events. The analog next-event estimator shows an interesting feature: it is the only one with a dual effect when increasing $\Sigma_a L$. It first increases the variance, but later decreases it. This decrease cannot be explained by the increased amount of scoring or reweighing events, which does not happen with this estimator. Instead, cutting of the simulation prematurely is what is responsible for the decreased statistical error. As also explained in the context of the non-analog track-length type next-event estimator (III), there is no variance if no scattering occurs. This tends also to be true if the survival probability is very low. For intermediate values, this is no longer true and early and late absorptions make up another important source of variance.

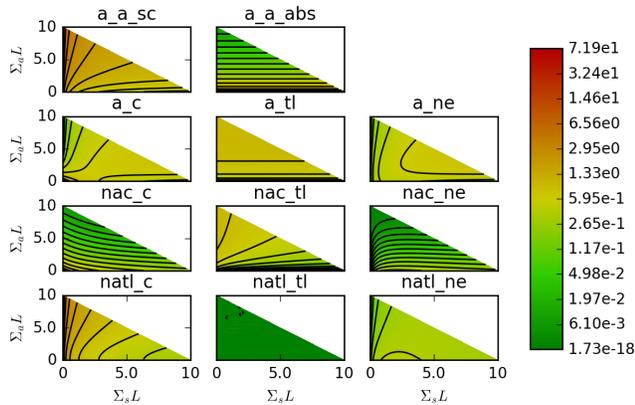


Fig. 2: The relative standard deviation for the different simulation-estimator combinations with $P_r = 1$.

B. Non-trivial scattering

Most of the effects at play already surfaced in the situation in which the scattering was trivial. There are a few new effects however, when the particle can change direction upon scattering. We will introduce this in Fig. 3 where the probability of going right after a scattering event equals 75%.

Again, the analog absorption estimator (III) provides a good starting point. In contrast to the situation with delta scattering, scattering collisions do influence the path of the particle now. This is a new source of stochasticity and increases the variance of the score. The negative effect of non-trivial scattering is visible in every subfigure, as can be seen by comparing Fig. 3 with Fig. 2.

The new negative aspect of an increase in the scattering rate combines with the already existing effects when there was only trivial scattering and results in a more complex picture for most estimators. Notably the collision estimators (III), the analog absorption estimator (III) and the non-analog collision type track-length estimator (III) have lost their monotonous decrease of the relative standard deviation with an increase in $\Sigma_s L$. For each of these, there are values of $\Sigma_a L$ to be found for which there is first a decrease of the variance and then an increase. The reason lies in the ever-increasing complexity of the path, which after a certain time dominates the initial

decrease in variance due to extra scoring moments.

Other values of P_r than 1 or 0.75 do not show different effects, but the specific trade-off between the different effects at play changes. Different values are shown in Fig. 4, 5 and 6.

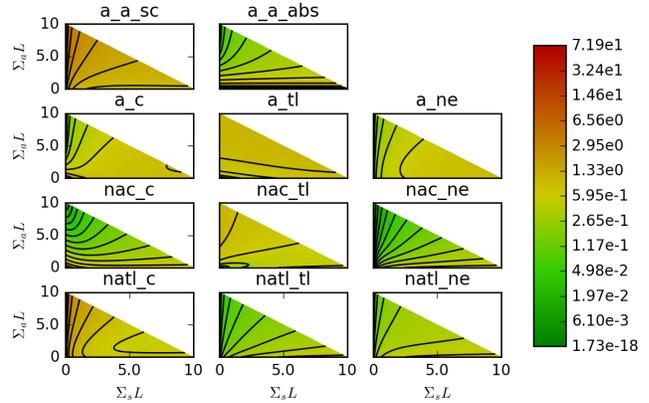


Fig. 3: The relative standard deviation for the different simulation-estimator combinations with $P_r = 0.75$.

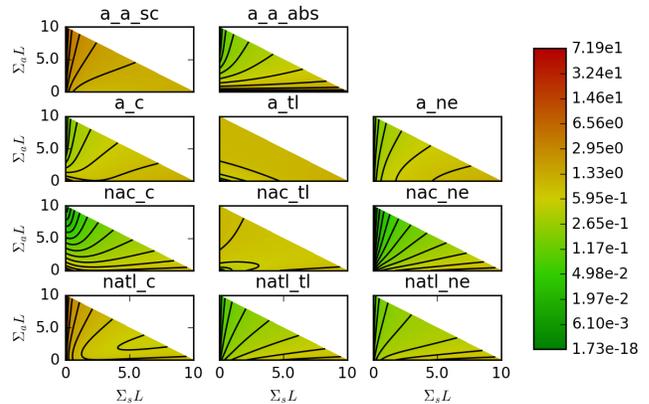


Fig. 4: The relative standard deviation for the different simulation-estimator combinations with $P_r = 0.5$.

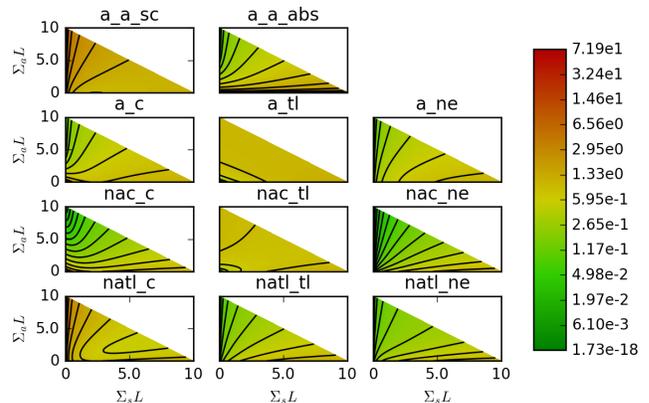


Fig. 5: The relative standard deviation for the different simulation-estimator combinations with $P_r = 0.25$.

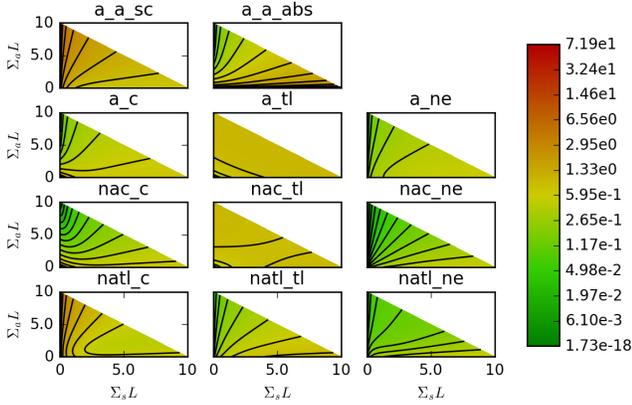


Fig. 6: The relative standard deviation for the different simulation-estimator combinations with $P_r = 0$.

2. Computational cost

The time needed to perform a Monte Carlo simulation of the neutral particle kinetic equation is typically proportional to the product of the number of collisions and the number of particles involved. Assuming the Central Limit Theorem holds, the relative error in a statistical sense is proportional to

$$\frac{\sigma}{\sqrt{N}},$$

with σ the standard deviation of the contribution of a particle. This means one can write the proportionality of the computational cost as

$$\text{cost} \propto \sigma^2 \mathbb{E}[\text{collisions per path}]. \quad (2)$$

In the same way as how ODEs for the variance were derived, formulas for the expected number of collisions can be found [17]. This is independent of the chosen estimator (a_abs, a_sc, c, tl, ne), but depends on the simulation type (a, nac, natl).

For the analog simulation (a), only the physical scattering collision, determined by the cross-section Σ_s have to be simulated as such. The absorption collision cuts the path and keeps the computational cost down this way. The non-analog track-length type simulations (natl) execute scattering collisions the same way as in analog simulations. Absorptions now no longer cut-off the particle path, but are only visible in the weight of the particle. In the non-analog collision type simulation (nac) all collisions are simulated as if they are scattering collisions. Hence, the number of scattering collisions that have to be simulated is the lowest for the analog simulations and the highest for the non-analog collision type simulations.

These new aspects complicate the figures even more and the results are shown for $P_r \in \{1, 0.5, 0\}$ in Fig. 7, 8 and 9.

Note that the ever increasing cost for the non-analog simulations with increasing $\Sigma_t L$ can in practice be avoided by using techniques like Russian roulette or weight cut-off [18].

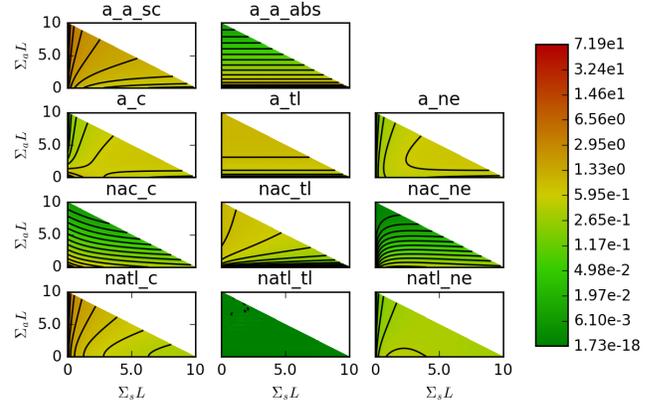


Fig. 7: A measure of computational cost for the different simulation-estimator combinations with $P_r = 1$.

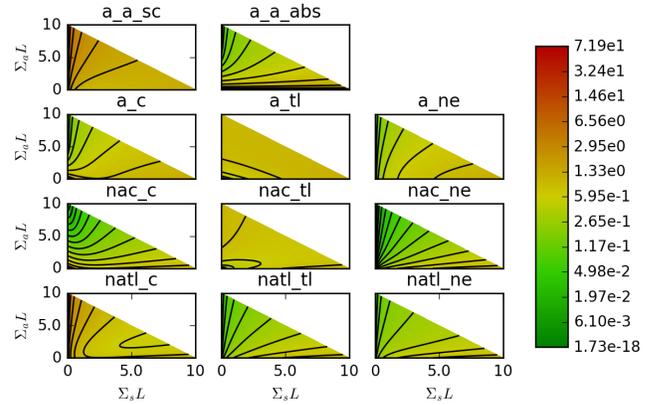


Fig. 8: A measure of computational cost for the different simulation-estimator combinations with $P_r = 0.5$.

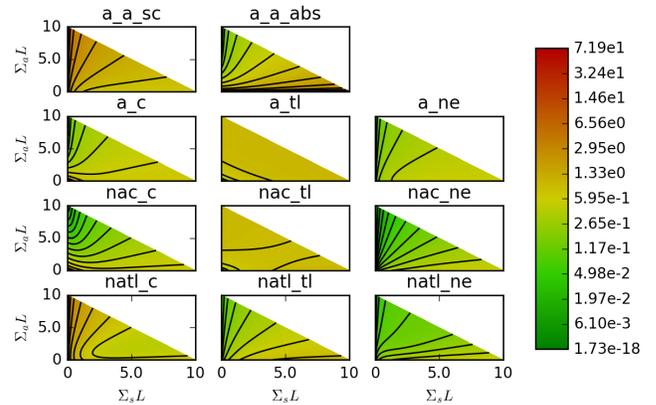


Fig. 9: A measure of computational cost for the different simulation-estimator combinations with $P_r = 0$.

3. Best estimator

An important question for practical use is which estimator to use for a given problem. From the previous pictures it was already possible to extract that there is not one estimator that

outperforms all others throughout the parameter domain. In this section we will make this more explicit, by splitting the parameter space depending on which estimator performs best. This will show that three estimators are competitive when regarding the relative standard deviation and four when taking the cost as defined in equation (2).

In the context of an actual simulation Σ_s/Σ_t and P_r can be seen as constant for a grid cell, even in the general 3D/3D simulation as discussed in section II. In contrast, $\Sigma_t L$ does depend on the particle. $\Sigma_t L$ is namely the inverse of the dimensionless free path length, so it is proportional to $1/|\mathbf{v}|$, with \mathbf{v} the velocity of the particle. By choosing $\Sigma_t L$ and Σ_s/Σ_t as axes in our figures, we only have to look at one vertical line when considering a single grid cell. Depending on the velocity of the particle, the relevant point of that line can be determined.

Fig. 10 shows for each point in the parameter domain the estimator with the lowest relative standard deviation, which corresponds to the lowest statistical error given a number of simulated particles. The pixels that are colored black, indicate that the numerical error in the simulation of the ODEs was larger than the difference between the computed variances. The graph does not show the full parameter space, as $\Sigma_t L$ can be unbounded, but it shows all the competitive estimation procedures and which parts of the domain they dominate. Fig. 10 shows the prevalence of the next-event estimator except for when there is a large amount of scattering. Then the track-length simulation-estimator combination becomes better. Towards $P_r = 1$ the natl_tl estimator becomes zero-variance, hence at $P_r = 1$, such a plot as in Fig. 10 would be cyan throughout.

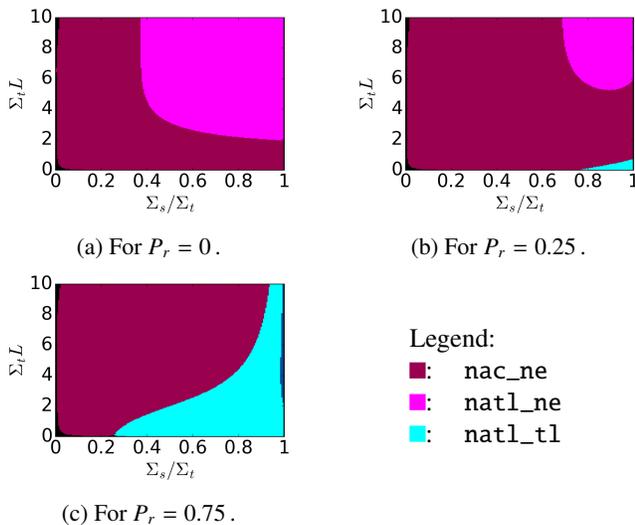


Fig. 10: The estimator with the lowest statistical error for a given number of simulated particles throughout the parameter space.

Fig. 11 shows the same for the computational cost. Note the connection with Fig. 10, of which the underlying value is squared and multiplied by the expected number of collisions, which only depends on the simulation type. The number of collisions is always lowest for a and highest for nac.

These effects can be seen when comparing of Fig. 11 and 10: natl_tl and natl_ex take over parts of the domain that went to nac_ex in Fig. 10. On top of that, there is a new estimation procedure that comes to the fore: a_ex. For small values of the survival fraction and not too low values of $\Sigma_t L$, non-analogous particles can remain in the domain for long without having a significant contribution to the estimation due to their low weight. Those constitute a computational cost, but have nearly no impact on the computed value. This is the reason analog simulations can also be the best option.

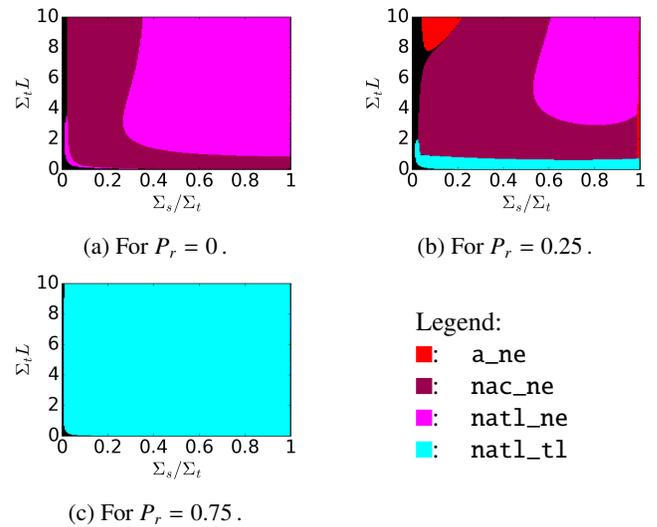


Fig. 11: The estimator with the lowest computational cost for a given statistical error throughout the parameter space.

Another question is what the gain in terms of statistical error and computational efficiency would be if you would not stick to only one estimator, but leave it free. We look at this for five values of P_r separately (0, 0.25, 0.5, 0.75 and 1), but for all values of Σ_s/Σ_t and $\Sigma_t L$. This is possible even though $\Sigma_t L$ can be any positive value because the standard deviation tends to a constant value for large $\Sigma_t L$. For the computational cost this is not true for non-analogous simulations (nac and natl), since the number of collisions becomes unbounded for $P_r \geq 0.5$, so then obviously only an analog estimation procedure can have the lowest value. The exception here is for $P_r = 1.0$ since the variance of natl equals 0 throughout the domain. Table I shows which estimation procedure has the lowest factor with which one might improve the relative statistical error or computational efficiency throughout the parameter domain, together with this maximal value.

Based on the values in the table, it is clear that there is no one-size-fits-all estimation procedure, even when looking at the different values of P_r separately. For the cases with trivial scattering ($P_r = 1$), there is no potential gain to be made, since natl_tl has zero statistical error throughout the domain.

P_r	Statistical error		Cost	
	Estimator	Maximal gain	Estimator	Maximal gain
0	nac_ne	1.4	nac_ne	1.9
0.25	nac_ne	1.2	a_ne	1.5
0.5	nac_ne	1.4	a_ne	2.2
0.75	nac_ne	2.0	a_ne	4.7
1	natl_tl	/	natl_tl	/

TABLE I: The single best estimation procedure choice for given P_r , together with the corresponding maximal factor of improvement if other estimation procedures would be allowed as well.

VII. DISCUSSION

Our results apply to a 1D0D simulation with homogeneous values for $\Sigma_t L$, Σ_s/Σ_t and P_r . To apply to a 1D0D our results of which estimation procedure performs best is not efficient since this computation is more complex than computing the result itself. Our results can however be used as a prediction of which estimator performs best in similar situations. This way the potentially large gain of table I of the previous section might be benefited from.

The aim of future research is to extrapolate these results to the MC subsimulation in B2-EIRENE simulations. Each of the grid cells there has homogeneous values for the scattering rate, survival probability and post-collision velocity distribution, and is as such quite similar to the situation here. The difference with the studied problem in this research lies in the higher dimensionality and the fact that the size of the velocity can change within the grid cell. The idea would be to select the estimation procedure depending on the parameters of a specific grid cell, or even to optimize the estimation procedure in this way by selecting an optimal linear combination of estimators. Formulas for the optimal weights — in a 1D0D scenario — can be found via invariant imbedding procedures as well. This should lead to more efficient calculations if there is a large difference in parameter values for the different grid cells, which is typically true in fusion simulations.

VIII. CONCLUSION AND FURTHER PERSPECTIVES

We combined four types of estimators and three types of simulations into twenty estimation procedures for reaction rate estimation. We compared them based on two performance measures, for which eleven estimation procedures behave uniquely. First, we looked at the relative standard deviation which is proportional to the statistical error for a given set of particles that enter the simulation. The second performance measure is proportional to the computational cost for a given statistical error.

We derived ODEs for these two performance measures with an invariant imbedding method for a simplified 1D0D situation, that is nonetheless significantly more complex than in the literature. The degrees of freedom that are retained are the dimensionless total collision cross-section ($\Sigma_t L$), the survival probability (Σ_s/Σ_t) and the anisotropy, expressed by

the probability of having a positive post-collision velocity (P_r).

We used numerical simulations of the systems of ODEs to discuss which effects are responsible for increases and decreases in performance of the different estimation procedures. These different effects make it highly non-trivial to select the estimation procedure with the best performance measure. We discuss how these results can be used to optimally select a local source term estimation procedure, based on the local values of the cross-sections and anisotropy. We also computed the potential loss of performance when sticking to one global estimation procedure. This loss is potentially very high.

ACKNOWLEDGMENTS

This work has been carried out within the framework of the EUROfusion Consortium and has received funding from the Euratom research and training programme 2014-2018 under grant agreement No.633053. The views and opinions expressed herein do not necessarily reflect those of the European Commission. BM is funded by a PhD fellowship of the Research Foundation - Flanders.

REFERENCES

1. D. REITER, M. BAELMANS, and P. BÖRNER, “The EIRENE and B2-EIRENE codes,” *Fusion Science and Technology*, **47**, 2, 172–186 (2005).
2. H. KAHN, “Applications of Monte Carlo,” (1956).
3. D. MACMILLAN, “Comparison of statistical estimators for neutron Monte Carlo calculations,” *Nuclear Science and Engineering*, **26**, 3, 366–372 (1966).
4. R. INDIRA, “Analytical study of leakage estimators in Monte-Carlo simulation of radiation transport,” *Annals of Nuclear Energy*, **15**, 5, 261–269 (1988).
5. I. LUX, “Variance and efficiency in transport, Monte Carlo Report KFKI-1979-35,” Tech. rep., Central Research Institute for Physics, Budapest, Hungary (1979).
6. I. LUX, “Systematic study of some standard variance reduction techniques,” *Nuclear Science and Engineering*, **67**, 3, 317–325 (1978).
7. J. SPANIER, “An analytic approach to variance reduction,” *SIAM Journal on Applied Mathematics*, **18**, 1, 172–190 (1970).
8. P. SARKAR and M. PRASAD, “Prediction of statistical error and optimization of biased Monte Carlo transport calculations,” *Nuclear science and engineering*, **70**, 3, 243–261 (1979).
9. R. INDIRA and T. JOHN, “Variance reduction under exponential and scattering angle biasing: An analytic approach,” *Nuclear Science and Engineering*, **94**, 4, 323–336 (1986).
10. R. INDIRA, “Optimization of biasing parameters in particle transport simulations,” *Annals of Nuclear Energy*, **16**, 11, 571–576 (1989).
11. C. J. SOLOMON, A. SOOD, T. E. BOOTH, and J. K. SHULTIS, “Verification of the history-score moment equations for weight-window variance reduction,” in “Proceedings of the International Conference on Mathematics

and Computational Methods Applied to Nuclear Science and Engineering (M&C 2011), Rio de Janeiro, RJ, Brazil,” (2011).

12. I. LUX and L. KOBLINGER, *Monte Carlo particle transport methods: neutron and photon calculations*, vol. 102, Citeseer (1991).
13. J. SPANIER, “Two pairs of families of estimators for transport problems,” *SIAM Journal on Applied Mathematics*, **14**, 4, 702–713 (1966).
14. D. MACMILLAN, “Note on a paper by J. Spanier concerning Monte Carlo estimators,” *SIAM Journal on Applied Mathematics*, **15**, 2, 264–268 (1967).
15. S. CHANDRASEKHAR, *Radiative transfer*, Oxford University Press (1950).
16. R. BELLMAN and G. M. WING, *An introduction to invariant imbedding*, John Wiley & Sons Inc (1975).
17. B. MORTIER, M. BAELMANS, and G. SAMAEY, in preparation.
18. J. SPANIER and E. M. GELBARD, *Monte Carlo principles and neutron transport problems*, Addison-Wesley Publishing Company (1969).