

Exact weight cancellation in Monte Carlo eigenvalue transport problemsHunter Belanger^{*,†}, Davide Mancusi[‡], and Andrea Zoia[‡]*Université Paris-Saclay, CEA, Service d'études des réacteurs et de mathématiques appliquées, 91191, Gif-sur-Yvette, France*

(Received 26 March 2021; accepted 23 June 2021; published 16 July 2021)

Random walks are frequently used as a model for very diverse physical phenomena. The Monte Carlo method is a versatile tool for the study of the properties of systems modeled as random walks. Often, each walker is associated with a statistical weight, used in the estimation of observable quantities. Weights are typically assumed to be positive; nonetheless, some applications require the use of positive and negative weights or complex weights and often pose particular challenges with convergence. In this paper we examine such a case from the field of nuclear reactor physics, where the negative particle weights prevent the power iteration algorithm from converging on the sought fundamental eigenstate of the Boltzmann transport equation. We demonstrate how the use of weight cancellation allows convergence on the physical eigenstate. To this end, we develop a method to perform weight cancellation in an exact manner, in three spatial dimensions. The viability of this algorithm is then demonstrated on a reactor physics problem.

DOI: [10.1103/PhysRevE.104.015306](https://doi.org/10.1103/PhysRevE.104.015306)**I. INTRODUCTION**

Many physical processes can be represented by the random movement of particles, or walkers, through phase space. Such phenomena include radiation transport, propagation of active molecules in living bodies, or the spread of epidemics [1–4]. Quite often, it is desirable to compute the properties of such systems with the aid of Monte Carlo simulations. This is certainly the case in neutron transport, where Monte Carlo techniques are used to sample the random walk process and thus solve the Boltzmann neutron transport equation [5]. In this context, the Monte Carlo method may be preferred to (faster) deterministic methods because it requires very few (if any) approximations: Monte Carlo simulations are typically used to obtain reference solutions to which deterministic solutions may be compared.

The random walkers have an associated statistical weight, which is typically positive and real. This weight is used to estimate the observable quantities which are sought in the simulation. Certain forms of transport problems, however, require the use of positive and negative weights, or even complex weights. Some examples include diffusion quantum Monte Carlo [6], the solution of neutron noise equations in the frequency domain [7], the determination of the second harmonic of the Boltzmann eigenvalue equation [8], and the determination of transmittance in graphics rendering [9]. Monte Carlo simulations that use negative or complex weights are notoriously difficult, as they often do not converge to the desired solution [6,8], or have a very high variance in the observed quantities [7]. Most literature recognizes that applying weight cancellation, where walkers carrying positive and negative

weights may annihilate with one another, is highly beneficial for the solution of these problems [6–8,10,11].

Several methods for performing weight cancellation have been proposed. One method, stemming from neutron transport, achieves weight cancellation by pairing walkers individually (as opposed to pooling weights of walkers within a defined region) [8]. While in theory this algorithm is exact (i.e., it does not introduce any approximation), it is difficult to implement and has a quadratic computational complexity in the number of random walkers being simulated. Another linear, exact regional cancellation algorithm has been proposed [11], but is only valid in one-dimensional (1D) geometries. One alternative approach consists of using an approximate technique of averaging the weight of all walkers within the same region [7]. This is easily applied to three dimensions and has linear complexity, but does not provide an exact solution. Finally, it must be mentioned that not all random walker problems are posed in a manner which makes cancellation difficult. In the full configuration interaction quantum Monte Carlo technique, for example, cancellation is straightforward, as walkers explore a discrete state space, as opposed to a continuous state space; when two walkers of different signs land in the same state, they may immediately annihilate [12].

In this paper we review some of the existing weight cancellation methods and propose an exact 3D regional cancellation algorithm that has linear complexity in the number of walkers and that can be easily generalized to any number of dimensions. The development of this method has been motivated by certain neutron transport problems that require the use of negative walker weights in conjunction with the power iteration technique (for obtaining the dominant eigenstate) and that fail to converge to the correct solution without weight cancellation; a simplified example of such a system is presented in Sec. II. Subsequently, in Sec. III we develop a mathematical model to explain why the power iteration fails in the presence of negative weights and why weight cancellation can resolve

*hunter.belanger@cea.fr

†davide.mancusi@cea.fr

‡andrea.zoia@cea.fr

this problem. We demonstrate how the existing 1D regional cancellation algorithm allows successful convergence of our 1D power iteration problem in Sec. IV, and we develop our 3D version of the cancellation algorithm in Sec. V, testing it on a reactor physics benchmark in Sec. VI. Finally, we present our conclusions in Sec. VII.

II. POWER ITERATION DEBACLE

In neutron transport, a main quantity of interest is the fundamental eigenpair of the Boltzmann k -eigenvalue equation for neutrons, which we write here in its one-speed form, as a function of position \mathbf{r} and direction $\hat{\Omega}$:

$$\begin{aligned} \hat{\Omega} \cdot \nabla \varphi_k(\mathbf{r}, \hat{\Omega}) + \Sigma_t(\mathbf{r})\varphi_k(\mathbf{r}, \hat{\Omega}) \\ = \int_{4\pi} \Sigma_s(\mathbf{r}, \hat{\Omega}' \rightarrow \hat{\Omega})\varphi_k(\mathbf{r}, \hat{\Omega}')d\hat{\Omega}' \\ + \frac{\nu(\mathbf{r})\Sigma_f(\mathbf{r})}{4\pi k} \int_{4\pi} \varphi_k(\mathbf{r}, \hat{\Omega}')d\hat{\Omega}'. \end{aligned} \quad (1)$$

Here φ_k is the eigenfunction, k is the eigenvalue, Σ_t is the total cross section, Σ_s is the scattering cross section, Σ_f is the fission cross section, and ν is the average number of new particles produced per fission. While it does not explicitly appear in Eq. (1), there is an implicitly defined capture cross section $\Sigma_c = \Sigma_t - \Sigma_s - \Sigma_f$, which results in the death of neutrons. In general, Eq. (1) admits several eigenvalue-eigenfunction pairs; the fundamental (largest) k eigenvalue is the multiplication factor, which is typically written as k_{eff} . The corresponding eigenfunction φ_0 is known as the angular neutron flux and it represents the average number of neutrons crossing a unit surface area per unit time [5].

In order to obtain k_{eff} and φ_0 , a numerical technique known as power iteration is often employed [5,13]. Although Eq. (1) is not written in such a form, we may assume that it represents an eigenvalue problem of the type $Lv_0 = \lambda_0 v_0$, with v_0 the fundamental eigenstate and λ_0 the fundamental eigenvalue. In relation to Monte Carlo transport problems, the operator L can be interpreted as the propagation of particles through a system from one fission event to another. With power iteration, one may obtain the fundamental state v_0 from any state b such that the inner product $\langle b, v_0 \rangle \neq 0$, by repeated application of the operator L :

$$\lim_{n \rightarrow \infty} \frac{L^n b}{|L^n b|} = v_0. \quad (2)$$

Therefore, starting from almost any initial state, it is possible to converge to the fundamental mode [13].

We wish to use Monte Carlo to yield a solution to Eq. (1), in the sense that the average density of walkers in phase space should satisfy Eq. (1), and represent the fundamental eigenmode. In the context of Monte Carlo, the idea of the power iteration technique needs to be adapted as follows. Starting from an arbitrary set of walkers, we sample random walks for all walkers until their death. Along the walk, a walker may be randomly killed if the magnitude of its weight becomes too low or split if the magnitude becomes too large [14]. During the random walks of this first generation, new particles will be born from fission. These fission particles are stored in a bank and are attributed to the second generation.

Once the first generation has finished, the banked particles undergo population control, where those with small weights may randomly be killed, those with large weights may be split into multiple particles, and the net weight of all particles is normalized. This is done to keep the net weight of all the particles at the beginning of a generation constant and to keep the number of particles in the simulation constant on average [14]. This culled and normalized particle bank is subsequently used as the source for the second generation. These particles then undergo the random walk, producing fission particles which will belong to the third generation. This application of the random walk mechanics on a generation of particles may be continued indefinitely. After a number of generations, the positions (and directions) of the fission particles will settle on an equilibrium distribution, representing the converged fission source for the problem [14]. The combined random walks of all particles from this converged fission source represent the dominant eigenfunction of Eq. (1). With the fission source converged, observable quantities (such as the angular flux and the multiplication factor) may be estimated in each generation. With an estimation of the angular flux and the multiplication factor obtained by each generation, an average may be calculated, although the estimation of the uncertainty is not trivial [15].

Within a generation, the random walk process is sampled as follows. A particle begins with an initial position and direction, which for the first generation may be sampled from a somewhat arbitrary distribution. The first task is to sample the distance the particle will fly before having a collision. In most applications of Monte Carlo simulation to reactor physics problems, the total cross section is usually assumed to be piecewise constant within each macroscopic geometric region composing the modeled system. To sample a flight distance in a material region with a spatially constant cross section, a random variable $\xi \sim \mathcal{U}[0, 1]$ is drawn and the cumulative distribution function for the flight distance must be inverted to obtain the distance to collision d_c :¹

$$\xi = \int_0^{d_c} \Sigma_t \exp(-\Sigma_t s) ds \Rightarrow d_c = -\frac{\ln(1 - \xi)}{\Sigma_t}. \quad (3)$$

The sampled distance is only valid in the given material region. If the distance d_m to the next material boundary is less than d_c , the particle is only moved by d_m , the cross section is updated, and a new distance to collision is sampled. Once at a collision site, the collision mechanics may be simulated. The particle can be captured with probability Σ_c/Σ_t , whereupon the history is terminated. Alternatively, scattering events may occur with probability Σ_s/Σ_t or the particle may produce new fission particles with probability Σ_f/Σ_t , which will be part of the next generation [14]. The multiplication factor may be estimated in several different manners: here we use the collision estimator. When a particle of weight w has a collision, we add the expected number of fission neutrons produced per collision at that location, namely, $w\nu(\mathbf{r})\Sigma_f(\mathbf{r})/\Sigma_t(\mathbf{r})$, to an accumulator which we refer to as K . Once the generation

¹It is for this reason that the underlying random walk is called exponential flight [16].

has finished, the estimate for the multiplication factor for the generation G is

$$k_{\text{eff}}^{(G)} = \frac{K}{W}, \quad (4)$$

with G the generation number and W the sum of the weights of all the particles at the beginning of the generation (which is a constant, because of the normalization procedure described above). The accumulator K must of course be reset to zero before each new generation begins.

A more realistic description of the system would consist of relaxing the hypothesis of piecewise constant cross sections, especially in view of multiphysics problems [17], where the cross sections for neutron transport depend on complex space-dependent physical feedback mechanisms such as temperature and material density fields. Recently, we examined different methods of sampling the flight distance for the case of spatially varying cross sections [18]. Such methodologies are highly desirable for the next generation of Monte Carlo transport codes, as they could allow for a better representation of the system being simulated, reduce memory requirements, and pair well with the reactor physics community's goals of taking into account a multiphysics approach [19]. In principle, sampling the distance to collision for spatially dependent cross sections would require inverting the following equation for d_c :

$$\xi = \int_0^{d_c} \Sigma_t(s\hat{\Omega} + \mathbf{r}_0) \exp\left(\int_0^s \Sigma_t(u\hat{\Omega} + \mathbf{r}_0) du\right) ds. \quad (5)$$

In practice, this is quite difficult to accomplish. Our previous work considered several algorithms which could sample flight distances from spatially continuous cross sections, without needing to directly invert Eq. (5). In particular, we examined the traditional delta tracking [20,21] and the negative weighted delta tracking methods [22,23].

In order to illustrate how these algorithms work, we revisit a simple one-dimensional transport problem that was considered in our previous paper. In the so-called rod model, particles may only move in the forward or backward direction, along a line segment. For our purposes, this line segment has a finite length and it is possible for particles to leak out of either end of the line. For our application we will set the boundaries of the segment at $x = 0$ and $x = 2$. As an example of a space-dependent cross section, we will use the broad Gaussian cross section from our earlier paper [18]. This cross section profile was chosen as it is not monotonic and better represents the idea that particles may see both an increase and a decrease in the cross section along their ray of flight. The chosen $\Sigma_t(x)$ has the form

$$\Sigma_t(x) = \sqrt{\frac{2}{\pi}} e^{-(x-1.23)^2} + 0.1 \quad \forall x \in [0, 2]. \quad (6)$$

The scattering, absorption, and fission probabilities are chosen to be spatially constant, having values of $\Sigma_s/\Sigma_t = 0.7$, $\Sigma_c/\Sigma_t = 0.1$, and $\Sigma_f/\Sigma_t = 0.2$, respectively. The average number of neutrons born per fission is $\nu = 2.5$; both scattering and fission are isotropic (equal probability of emission forward or backward).

Transport is conducted using the negative weighted delta tracking (NWDT) method [22,23]. The basic procedures for

Algorithm 1. Negative weighted delta tracking.

1. Sample uniform random variable ξ_1 ;
 2. $d := -\ln(\xi_1)/\Sigma_{\text{smp}}$;
 3. $\mathbf{r}_1 := d\hat{\Omega} + \mathbf{r}_0$;
 4. Sample random variable ξ_2 ;
 5. **if** $\xi_2 < q$ **then**
 6. $w := w \frac{\Sigma_t(\mathbf{r}_1)}{q\Sigma_{\text{smp}}}$;
 7. Perform real collision;
 8. **else**
 9. $w := w \frac{1 - \frac{\Sigma_t(\mathbf{r}_1)}{\Sigma_{\text{smp}}}}{1 - q}$;
 10. Virtual collision; go to line 1;
 11. **end**
-

this method are presented in Algorithm 1. It relies on the concept of real and virtual collisions. Negative weighted delta tracking requires two parameters: a sampling cross section Σ_{smp} , which is used to sample the distance to a tentative collision site, and a probability q , which is used to determine whether a collision is real or virtual. At a real collision, the mechanics of a collision are used to change the direction of the particle accordingly. A virtual collision does not simulate the collision mechanics; instead, the direction of the particle is left unchanged and the distance to a new tentative collision site is sampled, where the process begins again. Both Σ_{smp} and q are allowed to be functions of the position. According to the NWDT algorithm, the particle weight changes sign when a virtual collision is sampled, and $\Sigma_t > \Sigma_{\text{smp}}$. Here we have chosen $\Sigma_{\text{smp}} = 0.85\Sigma_t(x = 1.23)$, which leads to particles changing sign within the region $0.8 \lesssim x \lesssim 1.66$. As for q , we have used $q = \Sigma_t/(\Sigma_t + |\Sigma_{\text{smp}} - \Sigma_t|)$, which corresponds to the strategy proposed by Carter *et al.* [23,24]. This choice of q was motivated by our previous work, where the method demonstrated a very reasonable performance compared to the standard of delta tracking [18].

Delta tracking is a special case of NWDT, where the sampling cross section Σ_{smp} is taken to be a strict majorant cross section Σ_{maj} , with $\Sigma_{\text{maj}} \geq \Sigma_t$, and the acceptance probability q is taken equal to $\Sigma_t/\Sigma_{\text{maj}}$ [20,21].² With these choices, particle weights never change sign; however, insisting on the use of delta tracking can raise a number of practical problems, because it could be very challenging to obtain a majorant when using spatially continuous cross sections in a more realistic context. If one were to conduct a simulation where both the isotopic density and the temperature of materials varied spatially, there is in general no way to exactly determine the majorant cross section for the system. It is certainly possible to probe the phase space of the problem, testing the cross section at each point. However, if the selected sampling cross section is not in fact a majorant, it will be impossible to handle collisions for which $\Sigma_t > \Sigma_{\text{smp}}$; even if such an event does not occur, the results will be slightly biased. Conversely, if Σ_{smp} is chosen by applying a large safety margin to the maximum known cross section (in an effort to ensure that

²Delta tracking is also sometimes referred to as self-scattering in the electron transport community [25], as thinning [26], or as the null-collision method [27].




FIG. 1. Number of positive and negative particles per generation when using the population control scheme provided in Eq. (7).

underestimation of the majorant does not occur), the method becomes quite inefficient, as many unnecessary virtual collisions will occur [21]. Negative weighted delta tracking tolerates underestimations in the majorant cross section and also allows one to avoid gross overestimations of the majorant which reduce efficiency.

With NWDT, not only may there be negatively weighted particles in the system, but the weights may also change in magnitude. These two traits can lead to an increase in variance and simulation time. To mitigate these effects, roulette is used on particles with $|w| < 0.6$, with the survival weight being $w = \pm 1$ (ensuring that the particle keeps its initial sign). Particles are also split if $|w| \geq 2$.

In our previous work, we neglected fission. In view of testing NWDT in the framework of power iteration for eigenvalue problems, fission has been added to our model. The number n of new fission neutrons generated at any collision is taken to be

$$n = \left\lfloor |w| \frac{\nu \Sigma_f}{\Sigma_t} \frac{1}{k_{\text{eff}}^{(G-1)}} + \xi \right\rfloor, \quad (7)$$

where $\xi \sim \mathcal{U}[0, 1)$ is a uniform random variable and $k_{\text{eff}}^{(G-1)}$ is the estimated value of k_{eff} for the previous generation of particles. Without dividing by $k_{\text{eff}}^{(G-1)}$, the number of particles in the simulation would either increase exponentially if the system is supercritical ($k_{\text{eff}} > 1$) or decrease exponentially if the system is subcritical ($k_{\text{eff}} < 1$). This algorithm closely follows standard methods applied in Monte Carlo codes [28,29], with slight modifications to accommodate fission with negative particles. All new fission particles are born with a weight $w = \pm 1$, keeping the sign of the weight of the particle inducing fission. Between particle generations, all particle weights are multiplied by a normalization coefficient to ensure that the net weight of the system (sum of all particle weights) is always a constant value.

Figure 1 shows the number of positive and negative particles per generation, when starting with 10^6 particles uniformly distributed within the rod. Also shown is the net number of particles $N_{\text{net}} = N_+ - N_-$ and the total number of particles




FIG. 2. Weight of positive and negative particles per generation when using the population control scheme provided in Eq. (7).

$N_{\text{tot}} = N_+ + N_-$. Similarly, Fig. 2 shows the total positive weight W_+ (the sum of the weights of all positive particles), the total negative weight W_- (the sum of the magnitude of the weights of all negative particles), the net weight $W_{\text{net}} = W_+ - W_-$, and the total weight $W_{\text{tot}} = W_+ + W_-$. The number of particles is somewhat less meaningful than the weight, as one can use methods such as weight combing [30] to adjust how many particles are followed in the simulation, but this says nothing about the magnitude of their weight. Weight is the more natural quantity to examine, as its behavior is not controlled by the number of particles. It is instructive, however, to consider the interplay between the weight and number of particles. While Eq. (7) keeps the net number of particles stable as expected, the total number of particles is left to increase exponentially. Eventually, this increase in the number of particles causes each generation to take a longer amount of time to process, and the memory requirements increase as well. At some point, the computer memory is overwhelmed and the simulation fails. In an attempt to remedy this problem, we modified Eq. (7) by replacing $k_{\text{eff}}^{(G-1)}$; using the value of k_{eff} for the previous generation keeps the net number of particles constant on average, so a new quantity k_{tot} is defined, representing the increase in the total weight W_{tot} of the simulation. The quantity k_{tot} may be estimated in a similar manner to the multiplication factor, accumulating the collision estimator $|w| \nu(\mathbf{r}) \Sigma_f(\mathbf{r}) / \Sigma_t(\mathbf{r})$. Using $k_{\text{tot}}^{(G-1)}$ in Eq. (7), we obtain

$$n = \left\lfloor |w| \frac{\nu \Sigma_f}{\Sigma_t} \frac{1}{k_{\text{tot}}^{(G-1)}} + \xi \right\rfloor. \quad (8)$$

The effects on the particle populations from using this normalization technique are presented in Figs. 3 and 4. While the total number of particles indeed remains constant on average, the net number of particles now decreases exponentially. The net weight of course remains constant because we are normalizing it at the end of the generation. With fewer net particles to maintain the entire initial weight of the system, the magnitude of the weight of the particles increases drastically. This is indicated by the total weight of the system continuing




FIG. 3. Number of positive and negative particles per generation when using the population control scheme provided in Eq. (8).

to increase exponentially in Fig. 4. This leads to nearly equal quantities of positive and negative particles, each with very large weight magnitudes, causing large fluctuations in the flux estimators and therefore a larger variance in scores.

Other population control mechanisms were tested as well, hoping that they might stabilize the system. One such method was particle combing [30], conducted between each generation of particles. While combing did reduce the total number of particles going into each generation, the weight magnitude of each particle was quite large. Once they entered the simulation, the particles were split, causing an abrupt increase in memory consumption. After several generations, the simulation is killed by the operating system due to the large memory consumption.

Although the rod model example is admittedly simple, similar results have been obtained when using the NWDT algorithm in combination with power iteration on more realistic reactor physics problems in three dimensions and with energy-dependent cross sections (see Sec. VI). It is clear that the power iteration with NWDT cannot be used to estimate the




FIG. 4. Weight of positive and negative particles per generation when using the population control scheme provided in Eq. (8).

equilibrium distribution of walkers and therefore the solution to Eq. (1). It is therefore pertinent to better understand why the power iteration is failing and how the problem might be addressed.

III. NECESSITY OF WEIGHT CANCELLATION

A. Failure of power iteration

The Boltzmann equation presented in Eq. (1) is not capable of describing the interaction between the positive and negative particle populations in a simulation, only their combined average angular flux. It is fairly straightforward, however, to propose a formulation of the transport equation which is capable of representing both positively and negatively weighted particles. Considering Algorithm 1, outlining the NWDT method, we may write a set of coupled transport equations, one for the angular flux φ_+ of the positive particles and a second for the angular flux φ_- of the negative particles, resulting in the following coupled equations:

$$\hat{\Omega} \cdot \nabla \varphi_{\pm} + \Sigma_{\text{smp}} \varphi_{\pm} = S \varphi_{\pm} + \frac{1}{k} \mathcal{F} \varphi_{\pm} + \Delta(\Sigma_{\text{smp}} - \Sigma_t) \varphi_{\pm} + \Delta(\Sigma_t - \Sigma_{\text{smp}}) \varphi_{\mp}. \quad (9)$$

Here S is defined as

$$S \varphi = \int_{4\pi} \Sigma_s(\mathbf{r}, \hat{\Omega}' \rightarrow \hat{\Omega}) \varphi(\mathbf{r}, \hat{\Omega}') d\hat{\Omega}' \quad (10)$$

and \mathcal{F} is defined as

$$\mathcal{F} \varphi = \frac{v(\mathbf{r}) \Sigma_f(\mathbf{r})}{4\pi} \int_{4\pi} \varphi(\mathbf{r}, \hat{\Omega}') d\hat{\Omega}'. \quad (11)$$

We have also made use of the function

$$\Delta(x) = \begin{cases} x, & x > 0 \\ 0, & x \leq 0. \end{cases} \quad (12)$$

A rigorous derivation of this statement is given in the Appendix. There is a subtlety in this formulation: Instead of having one species of particles which may have a weight that is positive or negative, we now have two species (positive and negative) and both are represented by particles having a strictly positive weight. In this framework, our unknown is now the set of eigenfunctions

$$\zeta = \begin{bmatrix} \varphi_+ \\ \varphi_- \end{bmatrix}, \quad (13)$$

which is the solution of the generalized eigenvalue problem

$$A \zeta = \frac{1}{k} F \zeta, \quad (14)$$

where the operator A is

$$A = \begin{bmatrix} \mathcal{A}_{11} & \mathcal{A}_{12} \\ \mathcal{A}_{12} & \mathcal{A}_{11} \end{bmatrix}, \quad (15)$$

with

$$\begin{aligned} \mathcal{A}_{11} &= \hat{\Omega} \cdot \nabla + \Sigma_{\text{smp}} - S - \Delta(\Sigma_{\text{smp}} - \Sigma_t), \\ \mathcal{A}_{12} &= -\Delta(\Sigma_t - \Sigma_{\text{smp}}), \end{aligned} \quad (16)$$

and F is defined as

$$F = \begin{bmatrix} \mathcal{F} & 0 \\ 0 & \mathcal{F} \end{bmatrix}. \quad (17)$$

If we use \mathcal{V} to denote the vector space of the physical flux, then ξ is an element of the expanded vector space $\mathcal{V} \times \mathcal{V}$. Given the previous definitions, the physical flux is interpreted to be

$$\varphi = \varphi_+ - \varphi_- \quad (18)$$

and may be retrieved through the application of $\mathbf{D} : \mathcal{V} \times \mathcal{V} \rightarrow \mathcal{V}$, defined as

$$\mathbf{D} = \begin{bmatrix} I & -I \end{bmatrix} \quad (19)$$

and whose action is

$$\mathbf{D} \begin{bmatrix} f \\ g \end{bmatrix} = f - g. \quad (20)$$

The eigenvalue equation (14) has several interesting properties, which may be elegantly outlined by the introduction of the parity operator, namely,

$$\mathbf{P} = \begin{bmatrix} 0 & I \\ I & 0 \end{bmatrix}. \quad (21)$$

Equation (14) is invariant under the action of the parity operator; that is, if $\bar{\xi}$ is a solution of Eq. (14) with eigenvalue \bar{k} , then so is $\mathbf{P}\bar{\xi}$.

It is easily verifiable that $\mathbf{PAP} = \mathbf{A}$ and $\mathbf{PFP} = \mathbf{F}$ and therefore

$$\mathbf{A}(\mathbf{P}\bar{\xi}) = \frac{1}{\bar{k}} \mathbf{F}(\mathbf{P}\bar{\xi}). \quad (22)$$

Thus, if the eigenvalue \bar{k} is nondegenerate, we must have $\mathbf{P}\bar{\xi} = \pm \bar{\xi}$. This partitions the eigenstates into two sets, depending on the sign of their eigenvalue with respect to \mathbf{P} . The odd eigenstates have the form

$$\xi_o = \begin{bmatrix} \varphi \\ -\varphi \end{bmatrix}. \quad (23)$$

By inserting this ansatz into Eq. (14), we can verify that φ must solve the physical eigenvalue equation

$$\hat{\Omega} \cdot \nabla \varphi + \Sigma_t \varphi = S\varphi + \frac{1}{k} \mathcal{F}\varphi, \quad (24)$$

which is just Eq. (1). We denote the eigenfunctions and eigenvalues of this equation by φ_i and $k_{\varphi,i}$, respectively, with $i \in \{0, 1, \dots\}$. The even eigenstates have the form

$$\xi_e = \begin{bmatrix} \eta \\ \eta \end{bmatrix}, \quad (25)$$

where η must solve the modified Boltzmann equation

$$\hat{\Omega} \cdot \nabla \eta + \Sigma_{t,\eta} \eta = S\eta + \frac{1}{k} \mathcal{F}\eta, \quad (26)$$

which is verified by substitution into Eq. (14). Here we have defined

$$\Sigma_{t,\eta} = \Sigma_{\text{smp}} - |\Sigma_{\text{smp}} - \Sigma_t|. \quad (27)$$

We denote the eigenfunctions and eigenvalues of Eq. (26) by η_i and $k_{\eta,i}$, respectively, with $i \in \{0, 1, \dots\}$.

Since we always have $\Sigma_{t,\eta} \leq \Sigma_t$, the equation for η has the same scattering and fission terms as Eq. (1), but capture has been decreased, namely,

$$\Sigma_{c,\eta} = \Sigma_{t,\eta} - \Sigma_s - \Sigma_f \leq \Sigma_t - \Sigma_s - \Sigma_f. \quad (28)$$

Thus, on physical grounds, we correspondingly expect a larger dominant eigenvalue

$$k_{\eta,0} \geq k_{\varphi,0} = k_{\text{eff}}. \quad (29)$$

Therefore, the dominant eigenvalue of Eq. (14) is not k_{eff} , but the nonphysical eigenvalue $k_{\eta,0}$. Applying the power iteration method to Eq. (14) will result in convergence towards the latter eigenstate. Additionally, the equilibrium distribution is even under the exchange of positive and negative particles, i.e., it contains the same number of positive and negative particles and thus zero net particles. Converging to a state which has zero net particles is incompatible with traditional population control mechanisms (such as combing, or normalizing by k_{eff}), which are designed to keep the net weight constant, leading to the divergence in particle populations which was observed in Sec. II.

It is worth stressing the similarity between our analysis and the study by Spencer *et al.* [12] on the origin of the sign problem in the full configuration interaction quantum Monte Carlo without weight cancellation.

B. Modeling weight cancellation

Previous investigations have shown that weight cancellation can be very effective in dealing with Monte Carlo problems using particles of both positive and negative weights [6,8,12,31]. In view of these considerations, in the following we will formally address the effect of cancellation on Eq. (14). Weight cancellation can be modeled in the following manner. We start with some function of the phase space variables $(\mathbf{r}, \hat{\Omega})$, i.e., an element of \mathcal{V} . This function may be embedded in $\mathcal{V} \times \mathcal{V}$ by applying a suitable mapping $\mathbf{E} : \mathcal{V} \rightarrow \mathcal{V} \times \mathcal{V}$. There is some latitude in the definition of \mathbf{E} ; one possible definition is

$$\mathbf{E}_1 f = \begin{bmatrix} f \\ 0 \end{bmatrix}, \quad (30)$$

but equally valid choices could be

$$\mathbf{E}_0 f = \begin{bmatrix} 0 \\ -f \end{bmatrix}, \quad (31)$$

$$\mathbf{E}_{1/2} f = \begin{bmatrix} f/2 \\ -f/2 \end{bmatrix}, \quad (32)$$

or the nonlinear variant

$$\tilde{\mathbf{E}} f = \begin{bmatrix} \max(f, 0) \\ -\min(f, 0) \end{bmatrix}. \quad (33)$$

The only property that we require of \mathbf{E} is that it should be right inverse to \mathbf{D} , viz.,

$$\mathbf{DE} = \mathbf{I}. \quad (34)$$

This property expresses the fact that lifting a function from \mathcal{V} into $\mathcal{V} \times \mathcal{V}$ (the action of \mathbf{E}) followed by collapsing back into the space \mathcal{V} of physical fluxes (the action of \mathbf{D}) should not change the function we started with. It is easily verifiable that the previously proposed definitions of \mathbf{E} satisfy this property. In matrix notation, the operators \mathbf{E}_1 , \mathbf{E}_0 , and $\mathbf{E}_{1/2}$ are given by the general formula

$$\mathbf{E}_z = \begin{bmatrix} zI \\ (z-1)I \end{bmatrix}. \quad (35)$$

It is not possible to express \tilde{E} as a matrix, as it is nonlinear.

Let us now consider the operator product with the opposite ordering $C = ED$. This operator $C : \mathcal{V} \times \mathcal{V} \rightarrow \mathcal{V} \times \mathcal{V}$ is a projector, by virtue of Eq. (34):

$$C^2 = EDED = E(DE)D = ED = C. \quad (36)$$

By construction, the null space of C coincides with the null space of D , which is the space of even (unphysical) vectors:

$$C \begin{bmatrix} \eta \\ \eta \end{bmatrix} = 0. \quad (37)$$

Therefore, the operator C can be regarded as a model for perfect cancellation. We have shown that the even vectors represent the solutions of the nonphysical Boltzmann equation (26). This is supported by the explicit form of C which follows from the previous definitions for E :

$$C_z \begin{bmatrix} f_+ \\ f_- \end{bmatrix} = E_z D \begin{bmatrix} f_+ \\ f_- \end{bmatrix} = \begin{bmatrix} z(f_+ - f_-) \\ (1-z)(f_+ - f_-) \end{bmatrix}, \quad (38)$$

$$\tilde{C} \begin{bmatrix} f_+ \\ f_- \end{bmatrix} = \tilde{E} D \begin{bmatrix} f_+ \\ f_- \end{bmatrix} = \begin{bmatrix} \max(f_+ - f_-, 0) \\ -\min(f_+ - f_-, 0) \end{bmatrix}. \quad (39)$$

Note that only \tilde{C} can guarantee that both vector components are non-negative, but it requires that both C and E be nonlinear. In matrix notation, C_z reads

$$C_z = \begin{bmatrix} zI & -zI \\ (z-1)I & -(z-1)I \end{bmatrix}. \quad (40)$$

Returning to our eigenvalue equation, Eq. (14) may be modified to include the application of cancellation on the fission source:

$$A\zeta = \frac{1}{k}CF\zeta. \quad (41)$$

If ζ is even, $F\zeta$ must also be even as $[F, P] = 0$. Since C maps even vectors to 0, the cancellation operator in Eq. (41) causes the nonphysical eigenmodes corresponding to Eq. (26) to vanish.

At this point, it remains to be shown that cancellation does not perturb the physical eigenvalues of Eq. (14) that were initially associated with the odd eigenstates. To accomplish this, we start with Eq. (41) and apply D from the left on both sides, as this is the operator which performs the mapping $\mathcal{V} \times \mathcal{V} \rightarrow \mathcal{V}$:

$$DA\zeta = \frac{1}{k}DCF\zeta. \quad (42)$$

From Eq. (34) we may substitute $DC = D$. Using Eqs. (13), (15), and (18), this then simplifies to

$$(\mathcal{A}_{11} - \mathcal{A}_{12})\varphi = \frac{1}{k}\mathcal{F}\varphi. \quad (43)$$

Equation (12) indicates that $\Delta(x) - \Delta(-x) = x$, which, when combined with Eq. (16), yields

$$[\hat{\Omega} \cdot \nabla + \Sigma_t]\varphi = \mathcal{S}\varphi + \frac{1}{k}\mathcal{F}\varphi. \quad (44)$$

This is the physical Boltzmann transport equation. Therefore, cancellation leaves the physical eigenstates unchanged, so long as Eq. (34) holds true, regardless of the form of E .




FIG. 5. Behavior of the eigenvalue spectrum as a function of the ratio of the sampling cross section to the majorant cross section. Eigenvalues associated with the physical Boltzmann equation are shown by a solid line, while those associated with the nonphysical Boltzmann equation are shown by a dotted line.

In summary, the introduction of cancellation conserves all the (physical and nonphysical) eigenstates of Eq. (9); however, the nonphysical part of the spectrum is relegated to zero, while the physical eigenvalues are unperturbed and dominate the power iteration.

C. Deterministic proof of concept

In order to better understand the effects of cancellation and how it positively affects the convergence of the power iteration, it is fruitful to visualize the spectrum of the associated equations. For this purpose, we have written a deterministic solver that computes the full spectrum of eigenvalues and eigenfunctions (using the LAPACK library [32]) for a discretized version of the one-dimensional system introduced in Sec. II. A deterministic solver is necessary, since Monte Carlo can only be used to estimate the fundamental mode by power iteration. This solver uses a finite-difference method, with cross section values being taken at the midpoint of each spatial bin; both the positive and negative particle fluxes are explicitly considered.

Before considering cancellation, it is helpful to observe the behavior of the eigenvalue spectrum of Eq. (14) as a function of $\Sigma_{smp}/\Sigma_{maj}$. This quantity roughly measures the amount of negative weight that is produced by NWDT; the smaller the ratio is, the more negative weight is introduced. The spectrum is the union of the sets of eigenvalues $k_{\phi,i}$ and $k_{\eta,i}$ and is presented in Fig. 5. Naturally, the spectrum associated with the physical Boltzmann equation does not change with the ratio of $\Sigma_{smp}/\Sigma_{maj}$. The eigenvalues of the nonphysical Boltzmann equation diverge from the physical values as this ratio decreases from unity. The dominant nonphysical eigenvalue $k_{\eta,0}$ continually increases above $k_{\phi,0}$, while the other depicted eigenvalues decrease from their physical counterparts. In Eq. (29) we heuristically demonstrated that $k_{\eta,0} \geq k_{eff}$, and this is indeed the case in Fig. 5. Our heuristics, however, are not able to make any remarks as to the behavior of the higher




FIG. 6. Behavior of the eigenvalue spectrum as a function of cancellation efficiency α for the case of $\Sigma_{\text{smp}}/\Sigma_{\text{maj}} = 0.7$. Eigenvalues associated with the physical Boltzmann equation are shown by a solid line, while those associated with the nonphysical Boltzmann equation are shown by a dotted line.

eigenvalues, and while for this system $k_{\eta,i}$ for $i > 0$ always appears to decrease when $\Sigma_{\text{smp}}/\Sigma_{\text{maj}}$ decreases, we have observed other systems which do not exhibit this behavior.

When the cancellation operator \mathbf{C}_1 is added in accordance with Eqs. (40) and (41), we find that cancellation indeed suppresses the even eigenstates, leaving only the odd ones associated with the physical solutions, so long as $\Sigma_{\text{smp}}/\Sigma_{\text{maj}} > 0$. This is the case for a cancellation operator which is “perfect,” in that it is able to conduct cancellation in a manner which always completely neutralizes the negative particle population. While this is of course desired and possible to implement in a deterministic solver, we will later show in Sec. IV that 100% cancellation efficiency is not necessarily achievable in a Monte Carlo approach. To mimic this fact, we introduce an imperfect cancellation operator $\mathbf{C}^{(\alpha)}$:

$$\mathbf{C}^{(\alpha)} = \alpha \mathbf{C}_1 + (1 - \alpha) \mathbf{I}. \quad (45)$$

Here α represents the cancellation efficiency: When $\alpha = 1$ there is perfect cancellation and when $\alpha = 0$ there is no cancellation. This is a highly idealized approach to model imperfect cancellation and is not necessarily a faithful model of the effect of cancellation in a Monte Carlo setting.

Figure 6 presents the effects of cancellation with varying levels of efficiency for the case of $\Sigma_{\text{smp}}/\Sigma_{\text{maj}} = 0.7$. For values of α between 1 and approximately 0.21, the cancellation operator remains efficient enough that the physical eigenstate is the dominant one. A critical point is reached near $\alpha \approx 0.21$, where the eigenvalues of the physical and nonphysical systems are equal. Further reducing the efficiency of cancellation leads to the nonphysical eigenstate being the dominant one. This example indicates that for our system there is a *minimum amount of weight cancellation* required in order for power iteration to converge on the fundamental physical eigenstate. This is likely true for all systems which can be described by coupled transport equations for positive and negative particles.

Even in settings where enough cancellation is present to make the $k_{\phi,0}$ mode dominant, cancellation might still be unable to reduce $k_{\eta,0}$ to a level that allows for rapid statistical convergence towards the fundamental mode. To clarify this, we introduce the dominance ratio, which is defined as the ratio of the second-largest eigenvalue to the largest one. The dominance ratio is equal to $k_{\eta,0}/k_{\phi,0}$ for values of α where $k_{\phi,0} > k_{\eta,0} > k_{\phi,1}$. If instead α is such that $k_{\eta,0} < k_{\phi,1}$, then the dominance ratio is $k_{\phi,1}/k_{\phi,0}$, which is the dominance ratio of the physical system. The dominance ratio provides an indication of the rate of convergence of the power iteration. When it is very close to unity, more iterations are required to converge on the dominant eigenvalue. Should α be sufficiently large to ensure $k_{\eta,0} < k_{\phi,1}$, the convergence rate of the problem will no longer be bound by the efficiency of cancellation, but by the physical properties of the system being examined.

This model for cancellation in a deterministic context provides valuable insight into the behavior of this coupled system of positive and negative particles and the possible behavior of implementing cancellation of particle weights in a Monte Carlo context. It is for this reason that we do not go beyond the provided surface-level analysis of the effects of cancellation efficiency or choice of Σ_{smp} on the convergence of the deterministic model. We will now continue by discussing techniques of weight cancellation in Monte Carlo simulations.

IV. EXACT REGIONAL CANCELLATION SCHEME FOR 1D PROBLEMS

As mentioned in the Introduction, several cancellation strategies have been proposed in the past; one of these, devised by Booth and Gubernatis, is exact in 1D geometries [11]. This method works by partitioning all fissile domains of the problem domain into regions. Between each generation, weight cancellation among all new fission particles born in the region occurs. An important feature of this algorithm is that it has linear computational complexity with the number of particles partaking in cancellation, making it a good candidate for inclusion in general-purpose Monte Carlo transport codes. We have therefore chosen to focus on this algorithm. We will now provide a brief overview of how Booth and Gubernatis’ method works, as it is essential for understanding our proposed 3D algorithm, which will be developed in Sec. V.

Booth and Gubernatis make use of a quantity referred to as the fission density [11]. In an effort to reproduce and expand upon their work, we have chosen the following definition for the expected fission density: Considering a particle starting a flight at position x_0 and traveling in direction $\mu = \pm 1$, its expected fission density $f(x|x_0, \mu)$ is the expected value of the number of fission events per unit length around x . It can be written as the product of the probability density of flying from x_0 to x and having a collision at position x and the probability of that collision being in the fission reaction channel. The exact form of $f(x|x_0, \mu)$ depends on the transport methodology being employed. For NWDT, this formula can be deduced by examination of Algorithm 1. The probability density of flying from x_0 to x and having a real collision is

$$P_c(x|x_0, \mu) = \begin{cases} q \frac{\Sigma_t(x)}{q \Sigma_{\text{smp}}} \Sigma_{\text{smp}} e^{-\Sigma_{\text{smp}}|x-x_0|}, & \frac{x-x_0}{|x-x_0|} = \mu \\ 0, & \frac{x-x_0}{|x-x_0|} \neq \mu. \end{cases} \quad (46)$$




FIG. 7. Presented is a neutron (large colored circle), beginning a flight at x_0 and flying in the $\mu = 1$ direction (direction of flight is indicated by the arrow). For this flight, the expected fission density $f(x|x_0, \mu)$ is plotted as a blue line. The limits of the region are at x_{R-} and x_{R+} . The value of β displayed here is the minimum value of $f(x|x_0, \mu)$ within the region. Note that $f(x|x_0, \mu) = 0$ for positions that cannot be reached by the particle.

The $\Sigma_{\text{smp}} e^{-\Sigma_{\text{smp}}|x-x_0|}$ portion is the probability density of flying from x_0 to x and having either a real or a virtual collision. The factor q is the probability of the collision being real, while $\Sigma_t(x)/q\Sigma_{\text{smp}}$ is the weight correction factor for real collisions. This must of course be combined with the fission probability

$$P_f(x) = \frac{\Sigma_f(x)}{\Sigma_t(x)}. \quad (47)$$

As $f(x|x_0, \mu) = P_f(x)P_c(x|x_0, \mu)$, for the case of NWDT we may write

$$f(x|x_0, \mu) = \begin{cases} \Sigma_f(x)e^{-\Sigma_{\text{smp}}(x_0)|x-x_0|}, & \frac{x-x_0}{|x-x_0|} = \mu \\ 0, & \frac{x-x_0}{|x-x_0|} \neq \mu, \end{cases} \quad (48)$$

noting that the fission density is zero for all positions which cannot be reached by the particle during the considered flight.

With the expected fission density for a particle having been defined, let us consider a fission particle p belonging to the fission source and ready to start a random walk, at position x_p , located in an arbitrary region R (with bounds x_{R-} and x_{R+}). The parent of this particle (which was a member of the previous fission generation) was previously at position x_0 before flying to x_p and producing the current particle of interest. Figure 7 illustrates the expected fission density of the parent particle, as well as the bounds of region R and x_0 . The depicted value of β in the plot is the minimum value of the fission density within the region, for the flight of the parent particle:

$$\beta = \min_{x \in R} f(x|x_0, \mu). \quad (49)$$

The idea of Booth and Gubernatis's method is to estimate the fission density associated with the flight as the sum of a uniform component over the region R and a pointwise component concentrated at the actual fission site. For a fission particle born in R , which was induced by a particle beginning

a flight at x_0 and traveling in the direction μ , the bottom shaded portion in Fig. 7 represents the fraction of the fission density which is uniformly distributed in R , while the top portion represents the nonuniform fraction of the fission density, which depends on the position x_p where the parent lands and induces a fission. This indicates that, for the particle p , the fraction of the fission density concentrated at x_p is $[f(x_p|x_0, \mu) - \beta]/f(x_p|x_0, \mu)$ and the fraction uniformly distributed through region R is $\beta/f(x_p|x_0, \mu)$. As such, we can set the weight of p to be

$$w_p = w \frac{f(x_p|x_0, \mu) - \beta}{f(x_p|x_0, \mu)} \quad (50)$$

and at the same time create a uniform weight portion w_u of the particle, representing the portion of the particle which is evenly distributed throughout R , namely,

$$w_u = w \frac{\beta}{f(x_p|x_0, \mu)}. \quad (51)$$

It is important to note that $w_p + w_u = w$: The net weight in the system has not been modified; we have simply displaced a portion of w , distributing it uniformly through region R .

Dividing particles into a pointwise and uniform portion does not in and of itself accomplish any weight cancellation. When there are many fission particles within the same region, however, all of their uniform weight portions may be combined into a single weight which represents the uniformly distributed portion of the fission source for the region. This quantity is denoted by U_R and is the sum of all uniform weight portions for all particles born within the region, each coming with a sign. Negative particles will yield negative contributions to U_R , while positive particles will yield positive contributions, leading to a cancellation. Once all particles have contributed their uniform portion to U_R , n new particles are sampled uniformly within the region, where

$$n = \lceil |U_R| \rceil. \quad (52)$$

Each of these uniformly sampled particles within R then has a weight of U_R/n .

These newly sampled, uniformly distributed particles belong to the same generation as particle p and behave exactly like traditionally generated particles from this point on. Under this scheme, a particle which was initially positive will produce a positive uniform portion and keep a positive weight; conversely, negative particles will remain negative and produce negative uniform portions. In the case of x_0 being in the same region as x_p , the uniform weight portion w_u must be zero, as the minimum of the expected fission density "behind" the starting point of the flight is zero. Therefore, particles with x_0 and x_p in the same region effectively do not partake in the cancellation process. These two properties are ensured by β being the minimum value of f in the region. However, it is possible to relax the requirement that β should be the minimum of the expected fission density of the region. Booth and Gubernatis showed in their original work that any value of β may be used while still producing an unbiased result [11]. In the event of $\beta > f(x_p|x_0, \mu)$, the sign of the particle's pointwise weight portion will change.

With this regional method, the efficiency of cancellation can never be 100%. Only the uniform weight portions




FIG. 8. Positive weight W_+ , negative weight W_- , net weight W_{net} , and total weight W_{tot} in the 1D rod system obtained using 30 cancellation regions and 10^6 initial particles.

contribute to cancellation. For negative weighted fission particles, the pointwise weight portions do not cancel, always leaving some residual negative weight. Given this information, one might think that using a value of β which is larger than the minimum fission density might be beneficial, as it would increase the magnitude of the weight which goes into the uniform portion. This is not necessarily the case, however, as it could lead to the pointwise portions of positive particles becoming negative, via Eq. (50), possibly defeating the purpose of cancellation. Unfortunately, this method of cancellation cannot be modeled as a linear operator and therefore we are unable to put it in the context of the framework which was put forth in Sec. III. This makes regional cancellation conceptually similar to (though not the same as) the \tilde{C} operator mentioned in Eq. (39). It is inconsequential if our cancellation is linear or not; as long as it corresponds to a choice of E that preserves Eq. (34), the method will be exact and unbiased.

Results of the 1D methodology

We added the exact 1D cancellation method of Booth and Gubernatis to the power iteration problem outlined in Sec. II, performing the cancellation between generations and before the normalization of the system weight occurs. Thirty evenly spaced cancellation regions were used to partition the rod and β was always taken to be the minimum of the expected fission density within the region. With cancellation, the simulation was stable and was able to complete without issue. An eigenvalue of $k_{\text{eff}} = 0.33577 \pm 0.00005$ was obtained after 120 generations (20 inactive generations). This is in excellent agreement with the value obtained for the system using delta tracking, which was $k_{\text{eff}} = 0.33573 \pm 0.00005$ after the same number of generations. The weights W_+ , W_- , W_{net} , and W_{tot} are shown in Fig. 8. We see that the negative weight quickly rises to an equilibrium level, near 5000, and fluctuates then about that value for the duration of the simulation. The positive weight must naturally increase by the same amount to keep the net weight of the system constant. This leads to an increase in the total transported weight of approximately 10^4




FIG. 9. Number of positive (N_+), negative (N_-), net (N_{net}), and total (N_{tot}) particles in the 1D rod system obtained using 30 cancellation regions and 10^6 initial particles.

(the net weight still being 10^6). The evolution of the number of positive and negative particles is presented in Fig. 9. The number of particles in memory stabilizes near almost 2×10^6 , twice as many particles as were initially used.

The cancellation algorithm itself has linear computational complexity with respect to the number of fission particles at any given generation. However, it is less obvious whether the total number of fission particles present once the simulation has settled to equilibrium is itself a linear function of the net weight of the system. It is evident that the total amount of weight transported must increase until equilibrium is reached, but is the increase linear in the net weight of the system? To examine this, we perform several runs with different initial values of W_{net} and we look at the average value of W_-/W_{net} , as a function of W_{net} . The total weight can be deduced by looking at only the behavior of the negative weight, as the increase in the positive weight will mirror the negative weight. The result of this study is presented in Fig. 10. When too little net weight is injected in the system, a higher percentage of negative weight will be present in the system at equilibrium. There appears to be a critical point (in this case, near $W_{\text{net}} = 400$) above which the fraction of negative weight no longer decreases by adding more net weight. This might at first seem counterintuitive, but positive particles are always being converted to negative particles during the random walk (because of the NWDT algorithm) and the cancellation process is not 100% efficient. Above this critical point, no matter how much positive weight is added to the system at the beginning of the simulation, the equilibrium amount of negative weight will on average be a set fraction of the net weight. In other words, the total equilibrium weight is a linear function of the starting weight only for sufficiently large values of the latter.

Another consideration is the effect of the number of cancellation regions on the amount of negative weight. When there are fewer regions, each one must become larger; this makes cancellation less efficient, as β will decrease if one is taking β to be the minimum of f over the region. Conversely, making regions too small will result in too few particles which




FIG. 10. Negative weight fraction W_-/W_{net} as a function of net weight W_{net} in the 1D rod system obtained using cancellation with 30 cancellation regions.

can partake in cancellation, also reducing the efficiency. To examine this, we have plotted the negative weight fraction for the number of cancellation regions in Fig. 11. Indeed, from the figure, there is a range for which the fraction of negative weight decreases (i.e., the efficiency of cancellation increases) by adding more regions. The number of regions where the minimum occurs depends on the net weight of the system. Here this occurs near 6×10^2 regions for $W_{\text{net}} = 10^4$, 4×10^3 for $W_{\text{net}} = 10^5$, and 2×10^4 for $W_{\text{net}} = 10^6$, suggesting that the optimal number of regions does not quite scale linearly with the net initial weight. The three tested net weights all had very similar behavior and the efficiency improved at the same rate when adding cancellation regions, until the minimum value was reached. There is also a large range over which the negative weight fraction is less than 1% of the net weight. For all net weights, this range starts at approximately 20 regions and goes up to 2×10^3 regions for $W_{\text{net}} = 10^4$ or 3×10^5 regions for $W_{\text{net}} = 10^6$.

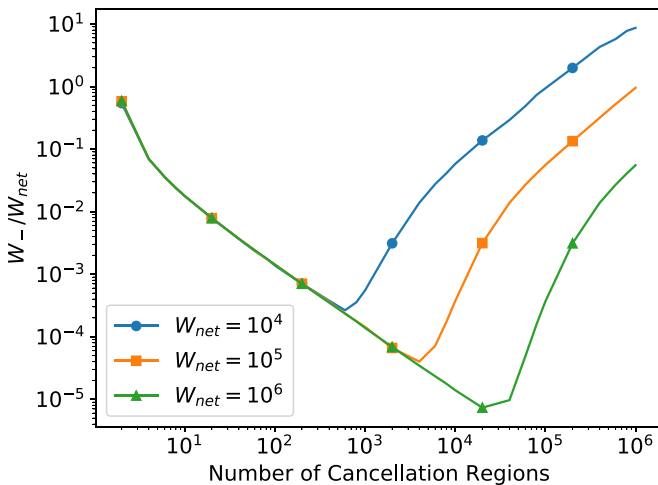


FIG. 11. Negative weight fraction W_-/W_{net} as a function of the number of cancellation regions for different net weights.

This section has outlined the methodology behind Booth and Gubernatis' regional cancellation. The potential of the method was successfully demonstrated, having canceled enough negative weight to allow for the proper convergence of the power iteration algorithm when using NWDT. In light of this success, we move on to extending the method to work with more realistic problems in higher spatial dimensions.

V. EXACT 3D, MULTIGROUP, REGIONAL CANCELLATION SCHEME

Unfortunately, it is not possible to directly implement the 1D regional cancellation scheme of Booth and Gubernatis in higher spatial dimensions. One-dimensional systems have the special property that lines coincide with volumes, ensuring that the lines of flight of all particles traversing the region overlap, as well as their uniform weight portions. This fact allows us to add the uniform weight portions of all particles in the region, effectively leading to a cancellation of negative and positive weights. For higher dimensions, this is not the case. The uniform portion of the particle weights may only be distributed along their rays of flight within the cancellation region, but the flight rays will never overlap completely. It is then no longer possible to combine the uniform weight portions as before by simply taking their sum.

The problem with the method of calculating the fission density in Eq. (48) is that it is only valid along the ray of flight of the parent particle, after the flight direction has been sampled. To extend the regional cancellation to higher dimensions, it is required to consider the probability of the parent particle scattering into the solid angle which is subtended by the cancellation region. This is done by examining $\hat{\Omega}'$, the direction of the parent particle at position \mathbf{r}_0 before the direction is changed by the scattering kernel. After the parent has been modified by the process of scattering, it then has a direction $\hat{\Omega}$, which must intersect the cancellation region (given that we are attempting to perform cancellation for a particle in the region which was induced by the parent). The scattering cosine for the interaction which the parent underwent at \mathbf{r}_0 is then $\mu = \hat{\Omega}' \cdot \hat{\Omega}$. For neutron transport, the azimuthal direction of a scatter is almost always isotropic, and only the scattering cosine is anisotropic (should there be any anisotropy) [5]. The probability density of scattering in direction $\hat{\Omega}$ from direction $\hat{\Omega}'$ is therefore

$$P(\hat{\Omega}|\hat{\Omega}') = \frac{P(\mu)}{2\pi}, \quad (53)$$

with $\int_{-1}^1 P(\mu) d\mu = 1$. Modifying Eq. (48) to use the flight kernel for three dimensions and applying this factor to consider the scattering angle, we arrive at the expected fission density

$$f(\mathbf{r}|\mathbf{r}_0, \hat{\Omega}') = \frac{P(\mu)\Sigma_f(\mathbf{r})}{2\pi|\mathbf{r} - \mathbf{r}_0|^2} e^{-\Sigma_{\text{smp}}(\mathbf{r}_0)|\mathbf{r} - \mathbf{r}_0|}, \quad (54)$$

where

$$\mu = \frac{\mathbf{r} - \mathbf{r}_0}{|\mathbf{r} - \mathbf{r}_0|} \cdot \hat{\Omega}'. \quad (55)$$

Taking a closer look at Eq. (54), one may recognize it as being what Lux and Koblinger refer to as the next-event estimator,




FIG. 12. Depiction of the 3D regional cancellation process. To find the uniform weight fraction of the fission particle at \mathbf{r} , β must first be determined. This is done by calculating the expected fission density at all eight corners of the cancellation region (though only two are depicted by the dashed arrows) from the parent particle's previous position and direction (\mathbf{r}_0 and $\hat{\Omega}'$) and taking the minimum value.

for the fission rate at \mathbf{r} given a collision at \mathbf{r}_0 (see [14], Sec. 6.IV.A). Note that Eq. (54) represents the expected fission density over the scattering and the free flight following it; this should be contrasted with Eq. (48), which represents the expected fission density over the following free flight *only*.

While there is only one natural shape for a cancellation region in one dimension, there is an infinite number of possible shapes in three dimensions which one could use. In this work we consider right rectangular prisms for our cancellation region. This is due to the fact that it is very simple to sample positions within a right rectangular prism uniformly, which is required in the cancellation processes. We will choose β to

be the minimum value of the expected fission density within the cancellation region. If $\Sigma_f(\mathbf{r})$ is spatially varying within the cancellation region and scattering is anisotropic, it is difficult to determine the true minimum; when $\Sigma_f(\mathbf{r})$ is homogeneous within the region and scattering is isotropic, however, one may determine the true minimum of the expected fission density for the flight from \mathbf{r}_0 to \mathbf{r} by evaluating $f(\mathbf{r}|\mathbf{r}_0, \hat{\Omega}')$ at all eight corners of the prism, as the minimum must occur at one of those eight points. Figure 12 provides a depiction of the cancellation process. Using the minimum as the value of β for the flight, it is possible to calculate a pointwise weight portion and a uniform weight portion of the fission particle, as was done before in Eqs. (50) and (51). Once the uniform portions of all fission particles in the region have been collected, one may take the net uniform weight of the region and sample new particles to add to the fission bank uniformly within the prism.

Until now, for simplicity, we have considered only single-speed transport, with all cross sections being energy independent. With no energy dependence, one does not need to consider how to sample the energy of the uniform particles. In all realistic applications in reactor physics, however, all cross sections are continuous functions of the neutron energy and the emission spectrum for the energy of fission neutrons depends on the collided nuclide and on the incoming energy of the neutron which induced the fission [5]. Under this assumption, it would be impossible to collect uniform portions of all the fission particles together. Without this action, no cancellation occurs. This fact poses a difficulty in applying weight cancellation to continuous-energy systems.

An approximation which is often used in the reactor physics community is the multigroup treatment of the energy variable in the phase space, by using cross sections which are piecewise constant in energy. These energy intervals in which the cross section is constant are referred to as groups. We denote the g th energy group by E_g , with $g \in \{1, 2, \dots, N_g\}$. By convention, the group corresponding with the highest energy is E_1 , while the group with the lowest energy is E_{N_g} [5]. The multigroup k -eigenvalue Boltzmann equation reads

$$\begin{aligned} \hat{\Omega} \cdot \nabla \varphi(\mathbf{r}, \hat{\Omega}, E_g) + \Sigma_t(\mathbf{r}, E_g) \varphi(\mathbf{r}, \hat{\Omega}, E_g) &= \sum_{g'=1}^{N_g} \int_{4\pi} \Sigma_s(\mathbf{r}, \hat{\Omega}' \rightarrow \hat{\Omega}, E_{g'} \rightarrow E_g) \varphi(\mathbf{r}, \hat{\Omega}', E_{g'}) d\hat{\Omega}' \\ &+ \frac{1}{4\pi k} \sum_{g'=1}^{N_g} \chi(E_{g'} \rightarrow E_g) \nu(\mathbf{r}, E_{g'}) \Sigma_f(\mathbf{r}, E_{g'}) \int_{4\pi} \varphi(\mathbf{r}, \hat{\Omega}', E_{g'}) d\hat{\Omega}'. \end{aligned} \quad (56)$$

The structure of Eq. (56) is formally that of a system of particles with N_g species, coupled with each other by means of scattering or fission.

Here the fission spectrum $\chi(E_{g'} \rightarrow E_g)$ is written to show a dependence of the fission particle energy on the parent particle's incident energy. This dependence on the incoming particle energy is relatively weak and is often ignored in the multigroup approximation [5]. In this case, we may simply write $\chi(E_g)$, moving it outside the sum in Eq. (56). Under this approximation, we may collect the uniform portions of

all fission neutrons in the region together; regardless of the energy of their parent, their energy spectra are all the same. However, one must change the definition of the expected fission density to take energy into account. Equation (54) is then modified to include the energy group E_g of the particle that induces fission:

$$\begin{aligned} f(\mathbf{r}|\mathbf{r}_0, \hat{\Omega}', E_g) &= \frac{P(\mu, E_g) \Sigma_f(\mathbf{r}, E_g)}{2\pi |\mathbf{r} - \mathbf{r}_0|^2} \\ &\times e^{-\Sigma_{\text{smp}}(\mathbf{r}_0, E_g) |\mathbf{r} - \mathbf{r}_0|}. \end{aligned} \quad (57)$$

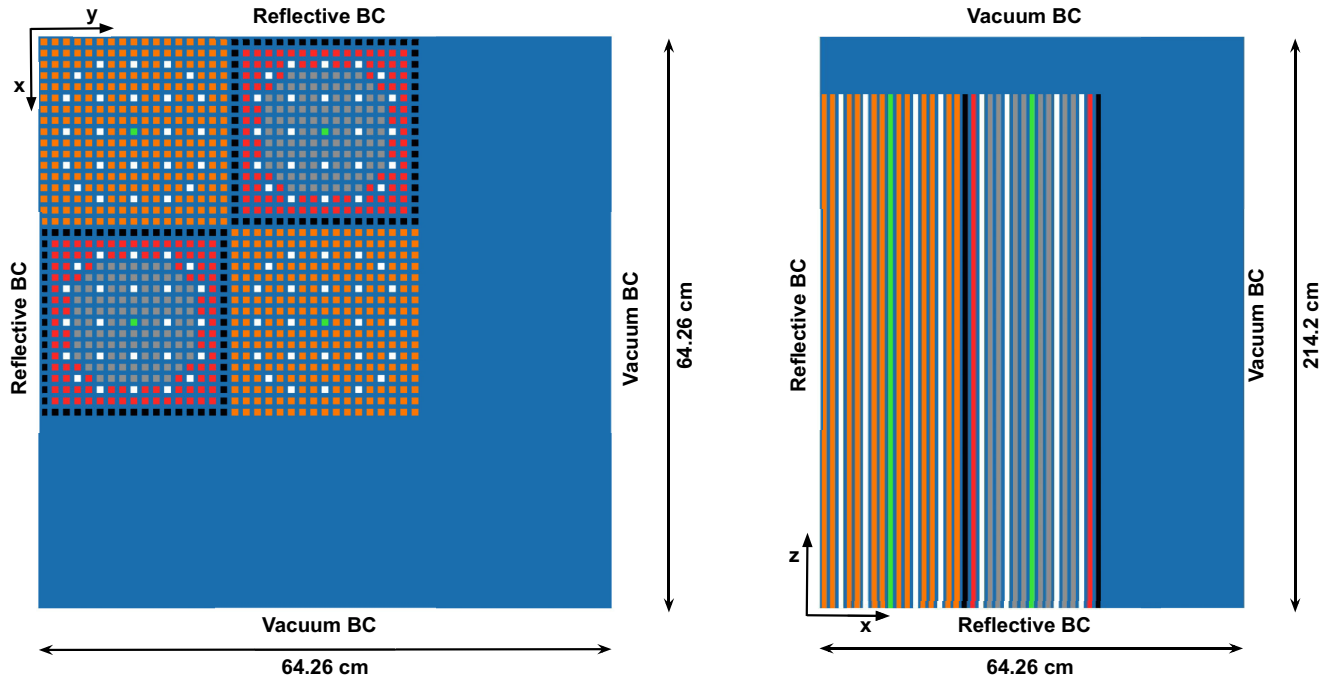


FIG. 13. Geometric configuration for our modified version of the C5G7 benchmark, with boundary conditions (BC). Each color represents a material associated with a unique set of cross sections.

A similar problem can occur for the direction of the fission particles as well. In continuous-energy transport, both the direction and energy of the fission particles may be a function of the energy and direction of the incoming neutron. While the nuclear data representations allow for this, it is a very marginal occurrence³ in major nuclear-data evaluation libraries [33,34]. Therefore, we do not need to make any special considerations for the angular distribution of fission particles in order to apply regional cancellation to multigroup problems.

We have seen that weight cancellation takes place at fission; however, one may wonder what is special about fission and whether one could perform cancellation at scattering events instead. At fission, the energy and angular distributions of the secondary fission particles are independent of the properties of the particle that induces the fission event. Fission represents the natural cancellation event because the three-dimensional distribution of the expected fission density [Eq. (57)] encodes all the six-dimensional distribution of the fission emission density in phase space.

VI. IMPLEMENTATION AND RESULTS

To test our exact 3D multigroup regional cancellation, we use a modified version of the C5G7 international benchmark, which we have depicted in Fig. 13 [35]. The C5G7 represents a small $\frac{1}{8}$ nuclear reactor core, with four fuel assemblies and seven energy groups, which is customarily used to assess and

compare deterministic transport codes. In its original specifications, the fuel pins within the assemblies are cylindrical, with a radius of 0.54 cm [35]. The use of cylindrical fuel pins however makes it difficult to use arbitrarily small cancellation regions while also being able to directly calculate the minimum value of Eq. (57), to be used as β . In order to facilitate cancellation, we have modified the benchmark so that the fuel pins are square in the x - y plane, with side lengths of 0.756 cm. This choice allows the fuel cells to be easily cut into an integral number of rectangular prisms to be used as cancellation regions, ensuring that no material other than fuel is present in the region. An important difference from the 1D rod system is that, in the C5G7 benchmark, the cross sections are piecewise spatially constant, being homogeneous within a given material cell. Also, while we derived the 3D cancellation formulas for the general case of anisotropic scattering, all scattering in the C5G7 benchmark is assumed to be isotropic, simplifying Eq. (57) to

$$f(\mathbf{r}|\mathbf{r}_0, E_g) = \frac{\Sigma_f(\mathbf{r}, E_g)}{4\pi|\mathbf{r} - \mathbf{r}_0|^2} e^{-\Sigma_{\text{smp}}(\mathbf{r}_0, E_g)|\mathbf{r} - \mathbf{r}_0|}. \quad (58)$$

Under this assumption, we have no need to keep track of the parent particle's initial direction $\hat{\Omega}'$, as there is equal probability of scattering in any direction.

The C5G7 benchmark makes use of three reflective boundary conditions, as shown in Fig. 13. One caveat of the regional cancellation method is that Eq. (57) is akin to a next-event point reaction rate estimator. Due to this, a trivial implementation of the method cannot be used with reflective boundary conditions [28]. Use of this method with reflective boundary conditions is only possible if special bookkeeping is done to

³One exception to this is the evaluation for ^{232}Th which has anisotropic distributions for prompt fission neutrons in the laboratory frame, in both ENDF/B VIII.0 and JEFF 3.3.

ensure a proper calculation of the expected fission density. If a particle begins at \mathbf{r}_0 and encounters a reflection before having a real collision at \mathbf{r} , then the distances which one would obtain by simply computing $|\mathbf{r} - \mathbf{r}_0|$ will not represent the actual distance traveled. It is also not enough to simply store the distance traveled by the particle, as one should ideally be able to calculate the value of f for any given \mathbf{r} in the region in order to calculate β . To ensure a proper calculation of $|\mathbf{r} - \mathbf{r}_0|$, the initial position \mathbf{r}_0 of the parent particle is stored: Upon each reflection, this position is transformed to the mirrored location on the other side of the plane of reflection. When doing this, the point will generally lie outside the defined geometry for the system, but this should not pose a problem, as we only need to be able to calculate the flight distance between \mathbf{r}_0 and other points in the cancellation region. With delta tracking and NWDT, $\Sigma_{\text{maj}}(E_g)$ and $\Sigma_{\text{smp}}(E_g)$ do not depend on the position within the geometry, but only on the energy group. We therefore do not need to worry about trying to find what material this fictitious point would be located in or the distance to the next surface.

A problem of this type is admittedly a simple case, where NWDT is not strictly necessary. One is able to trivially determine the majorant cross section so that delta tracking could be used, avoiding the problem of negative weights all together. As mentioned in Sec. II, however, in the context of multiphysics problems which are represented with spatially continuous cross sections, determining the majorant exactly is likely to be difficult, or impossible. It would certainly be interesting to test this algorithm on a more complex problem, where delta tracking would not be possible (or would be very inefficient). It would be difficult to verify that the presented cancellation method is working with such a problem, however, as it could not easily be solved with existing methods. This is why we have chosen to examine this simpler problem.

To test and evaluate cancellation and NWDT, a multi-group Monte Carlo code, called MGMC, was written to solve k -eigenvalue power iteration problems. MGMC supports general geometries using traditional surface-based descriptions of volumes. Either delta tracking [20] or the variant of negative weighted delta tracking developed by Carter *et al.* [22,24] may be selected for transport. Scalar flux and the fission reaction rate may be scored over a rectilinear mesh, using collision estimators.⁴ Monte Carlo estimates are saved as binary NumPy files [36] for easy analysis and plotting with PYTHON. Cancellation regions are defined by a rectilinear mesh imposed on top of the problem and can be used with both delta tracking and negative weighted delta tracking. Shared memory parallelism is implemented with OpenMP. Geometry, material properties, scores, cancellation, and simulation settings are all controlled with a YAML input file. MGMC is written in c++17 and has been made available under the CECILL v2.1 license [37]. One will also find the necessary input files to replicate our results there.

All simulations were started with 10^6 particles uniformly distributed across the four fuel assemblies, all in the first energy group, and ran for 2200 generations, with the first

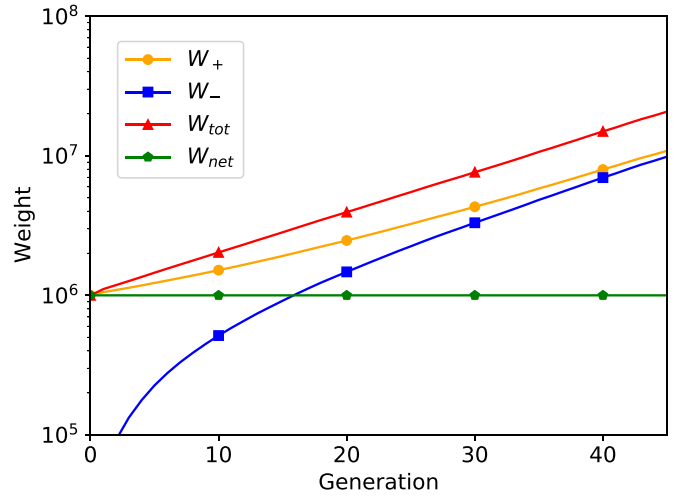


FIG. 14. Positive, negative, net, and total weights in the modified C5G7 benchmark obtained using negative weighted delta tracking and no weight cancellation.

200 generations being discarded to allow for fission source convergence. Delta tracking was first run (without weight cancellation) to obtain a reference multiplication factor for the system, which was found to be $k_{\text{eff}} = 1.219 \pm 0.00002$. When using the Carter *et al.* variant of negative weighted delta tracking, the majorant cross section was used for the sampling cross section in all the energy groups except for the first; for the first energy group, $0.9\Sigma_{\text{maj}}(E_1)$ was used as the sampling cross section. With this choice of Σ_{smp} , the total cross section is underestimated in all of the fuel pins in the first energy group. Any virtual collisions which occur in these regions of the phase space will therefore result in the particle weight changing sign.

Running this simulation with NWDT, and without the use of any weight cancellation, causes the particle populations to diverge and the simulation to fail. The exponential increase in total weight is depicted in Fig. 14. The behavior is nearly identical to the 1D case presented in Sec. II. Asymptotically, the negative weight will increase at the same rate as the positive weight, with the difference between the two remaining constant. The exponential increase in W_{tot} also leads to an exponential increase in the total number of particles in the simulation, overwhelming computer memory, exactly as in the 1D case. This supports the point made by Eq. (29); weight cancellation will always be necessary when attempting to perform power iteration simulations using NWDT.

In order to implement regional cancellation, all fuel pins were divided into cubical cancellation regions with side lengths of 0.252 cm. This mesh was chosen by trial and error, as we have found no practical way to know in advance for any system the minimum number of cancellation regions required to stabilize the simulation. Cubical cancellation regions also seemed to be more efficient than regions which had aspect ratios much larger than one, as compact region shapes led to a higher minimum value of f in the region and therefore increased β .

Our 3D multigroup exact regional cancellation algorithm was able to stabilize the particle populations when using

⁴Scalar flux is defined as $\int_{4\pi} \varphi(\mathbf{r}, \hat{\Omega}) d\hat{\Omega}$ [5].




FIG. 15. Scalar flux of the seventh energy group in the center axial slice of the core obtained using negative weighted delta tracking and 3D regional cancellation.

NWDT in conjunction with the previously outlined parameters and an eigenvalue of $k_{\text{eff}} = 1.219\,15 \pm 0.000\,05$ was obtained. This is in very good agreement with the value of k_{eff} obtained from traditional delta tracking with no weight cancellation, as the two estimates differ by less than one standard deviation. The thermal (seventh energy group) scalar flux which was tallied is shown in Fig. 15. This was compared with the flux tally from the delta tracking simulation by performing a Student t -test in each mesh cell [38]. The plot of the absolute value of the Student t -statistic is provided in Fig. 16. It can be seen there that the two flux estimates are also in good agreement with one another, with no apparent spatial dependence in their difference.

These two maps show only a small portion of the examined phase space. For a more thorough comparison, a histogram of the distribution of the Student t -statistic for the NWDT flux and the delta tracking flux is given in Fig. 17 [38]. Large portions of the flux tally mesh have average values of zero, due to the large water reflectors. These elements were removed before producing Fig. 17. Elements which had a relative error greater than 20% (in either the NWDT or the delta tracking score) were also removed, as the Student t -test is only applicable on normally distributed variables; therefore, we need to ensure that sufficient statistics are accumulated so that the central limit theorem may apply. The 20% cutoff is admittedly somewhat arbitrary, but it is likely that bins with such a high relative error are not normally distributed. It can be seen that the histogram has very good agreement with the theoretical distribution, which we assume to be normal given the large number of degrees of freedom for the comparison.

Finally, Fig. 18 shows the behavior of the weights in the system for the first 100 generations in the simulation. As with the 1D case, there is initially no negative weight. This




FIG. 16. Absolute value of the Student t -statistic for the difference between the delta tracking and NWDT flux estimates. The portion displayed is for the seventh energy group, at the center axial slice.

increases rapidly in the first 10–20 generations, before beginning to level out. By the time 80 generations have passed, the weights have reached their equilibrium values. This is much longer than the 1D case, which only required 2–3 generations before the weights reached equilibrium. The longer time to convergence is attributed to the C5G7 benchmark having a higher dominance ratio than the 1D case.

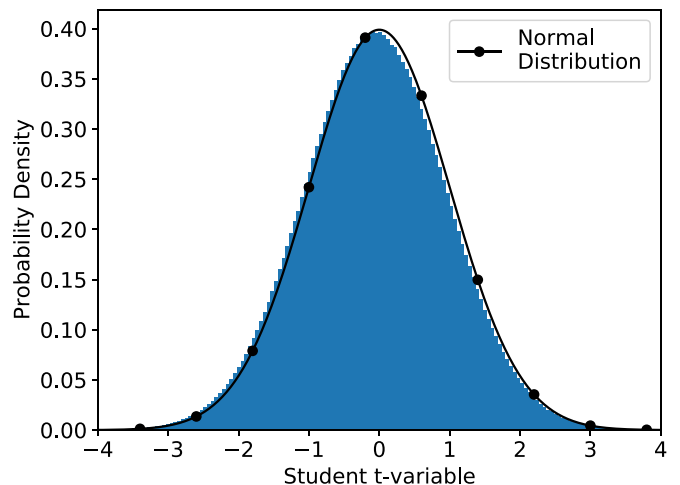


FIG. 17. Histogram of the Student t -statistic for the difference between the flux estimates obtained with delta tracking and NWDT. The curve plotted on top is the expected normal distribution.




FIG. 18. Positive, negative, net, and total weights in the modified C5G7 benchmark obtained using negative weighted delta tracking with exact 3D regional weight cancellation.

VII. CONCLUSION

Over the course of this work, we have presented a previously undocumented population control problem which arises when attempting to run k -eigenvalue power iteration Monte Carlo simulations with both negative and positive weights in negative weighted delta tracking [22]. Modeling the transport process through a set of coupled Boltzmann transport equations for negative and positive particles, we were able to show that the power iteration technique applied to negative weighted delta tracking will always fail to converge, as the sought physical fundamental eigenvalue is not the dominant eigenvalue of the system. Instead, a fictitious eigenvalue, associated with a system with a lower amount of absorption, is now dominant. We have formally developed a cancellation operator, demonstrating theoretically, and with a deterministic model, how particle weight cancellation can suppress the fictitious eigenstate and restore convergence of power iteration to the physical fundamental eigenvalue.

To demonstrate weight cancellation in a Monte Carlo context, we have implemented the exact 1D region cancellation algorithm of Booth and Gubernatis [11] in a 1D rod model. Weight cancellation did stabilize the particle populations through the fission generations, allowing the simulation to finish. As the method outlined by Booth and Gubernatis is only valid in 1D single-speed problems, we have developed an exact 3D regional cancellation method. We tested our 3D algorithm on a modified version of the C5G7 reactor physics benchmark, using negative weighted delta tracking. Cancellation stabilized the particle populations and also resulted in estimates for the multiplication factor and flux which were in agreement with the reference results obtained through delta tracking. Our exact 3D multigroup regional cancellation could potentially be useful for other applications in neutron transport as well. One such case is the use of Monte Carlo methods to obtain higher harmonics of the k -eigenvalue equation [8]. Our cancellation algorithm could also potentially prove useful in transport problems involving complex particle weights such

as the search for critical buckling or the solution of the neutron noise equations [7,39].

Several questions remain to be settled in regard to the methodology of 3D cancellation. First, in this work we have required that a cancellation region consist of only one material. Being able to have multiple materials in a cancellation region would make it easier to perform cancellation on different geometric forms (cylindrical pins). Second, extension of this algorithm to continuous energy will require the application of new techniques, because in continuous-energy transport the fission spectrum depends on the incoming neutron energy. Finally, the efficient choice of β should be evaluated. While we have chosen to use the minimum value of the expected fission density in the region for β , Booth and Gubernatis made clear that any value of β results in an unbiased estimate of the fission density. However, the variance of the weights of the fission particles clearly depends on β . It is desirable to characterize which choices of β result in more efficient cancellation. All of these issues will require further investigation to improve the algorithm and to probe its possible extension to continuous-energy problems.

APPENDIX: DERIVATION OF COUPLED BOLTZMANN TRANSPORT EQUATIONS

We provide here the derivation of the coupled system of Boltzmann equations, where the populations of positive and negative particles are treated separately and are assumed to be governed by the rules of negative weighted delta tracking (Algorithm 1) [22]. The quantities that we wish to describe are the angular fluxes of positive and negative particles. These quantities are not physical observables, but they are useful to characterize the behavior of the random walk. Consider any one of the usual Monte Carlo estimators for the (physical) angular flux; the unbiasedness condition requires any such estimator to be proportional to the particle weight w . Replacing the particle weight w in the estimator with

$$w_+ = \begin{cases} w, & w > 0 \\ 0, & w \leq 0 \end{cases} \quad (\text{A1})$$

results in a modified estimator; we define the angular flux of positive particles φ_+ to be the expected value of the modified estimator. Likewise, replacing the particle weight with

$$w_- = \begin{cases} 0, & w > 0 \\ -w, & w \leq 0 \end{cases} \quad (\text{A2})$$

results in a modified estimator, whose expected value is defined to be the angular flux of negative particles φ_- .

A few properties are worth stressing. First, if all particle weights are positive, then $\varphi_+ = \varphi$ (the physical angular flux) and $\varphi_- = 0$. Second, since $w_+ - w_- = w$, then $\varphi_+ - \varphi_- = \varphi$. Finally, since $w_+ \geq 0$ and $w_- \geq 0$, then both φ_+ and φ_- are always non-negative.

With these postulations, we may commence our derivation of the angular flux of positive particles. We will obey the rules for negative weighted delta tracking [22], presented in Algorithm 1. The Boltzmann equation for neutron transport is canonically written in the form of a balance equation, where losses equal gains. These losses and gains refer to the change in the neutron flux at the phase space point $(\mathbf{r}, \hat{\Omega})$. This can be

referenced in Eq. (1), where the left-hand side of the equation represents losses, while the right-hand side represents gains. More details on the reasoning and derivation behind particular terms can be found in [5].

Starting with the left-hand side, we first consider losses due to particle streaming L_S . This term is identical to that found in Eq. (1), only replacing the physical angular flux with the angular flux of positive particles:

$$L_S = \hat{\Omega} \cdot \nabla \varphi_+(\mathbf{r}, \hat{\Omega}). \quad (\text{A3})$$

Next we consider losses due to collisions L_C . In the case of NWDT, it is possible to be removed from the phase space point by having a collision at \mathbf{r} (real or virtual), which occurs with a cross section Σ_{smp} . We therefore consider the losses due to collisions with the term

$$L_C = \Sigma_{\text{smp}}(\mathbf{r})\varphi_+(\mathbf{r}, \hat{\Omega}). \quad (\text{A4})$$

In the event of a real collision, the particle direction will change, so there is indeed a removal from the phase space point. If a virtual collision occurs and the positive particle changes sign to be a negative particle, this is also a loss. If a virtual collision does not result in a positive particle becoming negative, then it is not actually a loss. This case will be treated on the right-hand side by adding a gain term for virtual collisions which do not result in a sign flip. The only possible sources for loss are then L_S and L_C , and we then have a left-hand side of

$$L_S + L_C = \hat{\Omega} \cdot \nabla \varphi_+(\mathbf{r}, \hat{\Omega}) + \Sigma_{\text{smp}}(\mathbf{r})\varphi_+(\mathbf{r}, \hat{\Omega}). \quad (\text{A5})$$

Next we may consider gains at the phase space point $(\mathbf{r}, \hat{\Omega})$. Gains may only come from particles which have had either a

real or a virtual collision, allowing them to enter the phase space point in question. First, we consider the gains from real collisions G_R . A real collision occurs with probability $q(\mathbf{r})$ and when this happens, a weight modification also occurs, which is a multiplication of the factor $\Sigma_t(\mathbf{r})/q(\mathbf{r})\Sigma_{\text{smp}}(\mathbf{r})$:

$$G_R = q(\mathbf{r}) \frac{\Sigma_t(\mathbf{r})}{q(\mathbf{r})\Sigma_{\text{smp}}(\mathbf{r})} R. \quad (\text{A6})$$

Here R is the sum of the gains of all real collision channels. We have two possible reaction channels which contribute to gains at $(\mathbf{r}, \hat{\Omega})$: particles having a collision at \mathbf{r} and scattering from direction $\hat{\Omega}'$ into direction $\hat{\Omega}$ (R_S) and particles having a collision at \mathbf{r} , inducing a fission particle which is born traveling in direction $\hat{\Omega}$ (R_F). Treating the scattering term, we must sum over all possible incoming directions, resulting in

$$R_S = \frac{\Sigma_{\text{smp}}(\mathbf{r})}{\Sigma_t(\mathbf{r})} \int_{4\pi} \Sigma_s(\mathbf{r}, \hat{\Omega}' \rightarrow \hat{\Omega}) \varphi_+(\mathbf{r}, \hat{\Omega}') d\hat{\Omega}'. \quad (\text{A7})$$

We must multiply by the ratio $\Sigma_{\text{smp}}/\Sigma_t$ to account for the fact that we are conditioning on the collision being real. We now perform a similar operation for the fission channel, with the assumption that the distribution of the direction of fission particles is isotropic:

$$R_F = \frac{\Sigma_{\text{smp}}(\mathbf{r})}{\Sigma_t(\mathbf{r})} \frac{v(\mathbf{r})}{4\pi k} \int_{4\pi} \Sigma_f(\mathbf{r}) \varphi_+(\mathbf{r}, \hat{\Omega}') d\hat{\Omega}'. \quad (\text{A8})$$

These being the only two real collision channels ($R = R_S + R_F$), we may combine the definitions of R_S , R_F , and G_R to obtain the result presented in the following equation, where we employ our previous definitions for \mathcal{S} and \mathcal{F} from Eqs. (10) and (11):

$$\begin{aligned} G_R &= q(\mathbf{r}) \frac{\Sigma_t(\mathbf{r})}{q(\mathbf{r})\Sigma_{\text{smp}}(\mathbf{r})} \left(\frac{\Sigma_{\text{smp}}(\mathbf{r})}{\Sigma_t(\mathbf{r})} \int_{4\pi} \Sigma_s(\mathbf{r}, \hat{\Omega}' \rightarrow \hat{\Omega}) \varphi_+(\mathbf{r}, \hat{\Omega}') d\hat{\Omega}' + \frac{\Sigma_{\text{smp}}(\mathbf{r})}{\Sigma_t(\mathbf{r})} \frac{v(\mathbf{r})}{4\pi k} \int_{4\pi} \Sigma_f(\mathbf{r}) \varphi_+(\mathbf{r}, \hat{\Omega}') d\hat{\Omega}' \right) \\ &= \int_{4\pi} \Sigma_s(\mathbf{r}, \hat{\Omega}' \rightarrow \hat{\Omega}) \varphi_+(\mathbf{r}, \hat{\Omega}') d\hat{\Omega}' + \frac{v(\mathbf{r})}{4\pi k} \int_{4\pi} \Sigma_f(\mathbf{r}) \varphi_+(\mathbf{r}, \hat{\Omega}') d\hat{\Omega}' \\ &= \mathcal{S}\varphi_+ + \frac{1}{k} \mathcal{F}\varphi_+. \end{aligned} \quad (\text{A9})$$

All that remains are gains from virtual collisions G_V . The probability of a virtual collision is $1 - q(\mathbf{r})$ and is accompanied by a weight modification of

$$\left| \frac{1 - \frac{\Sigma_t(\mathbf{r})}{\Sigma_{\text{smp}}(\mathbf{r})}}{1 - q(\mathbf{r})} \right|. \quad (\text{A10})$$

We take the absolute value here, as our particle weights may never become negative. Instead, a change in sign is modeled by the transfer of a particle from the positive population (corresponding to φ_+) to the negative population (corresponding to φ_-). Our virtual gains are

$$G_V = [1 - q(\mathbf{r})] \frac{\left| 1 - \frac{\Sigma_t(\mathbf{r})}{\Sigma_{\text{smp}}(\mathbf{r})} \right|}{1 - q(\mathbf{r})} V, \quad (\text{A11})$$

with V the sum of gains due to virtual collisions. The first channel for virtual collisions is the previously mentioned case

of a positive particle having a virtual collision and remaining positive (V_+). This only occurs when $\Sigma_{\text{smp}}(\mathbf{r}) \geq \Sigma_t(\mathbf{r})$ and can be modeled with the Heaviside function

$$\Theta(x) = \begin{cases} 1, & x \geq 0 \\ 0, & x < 0. \end{cases} \quad (\text{A12})$$

Our gains from virtual collisions of positive particles is then

$$V_+ = \Sigma_{\text{smp}}(\mathbf{r}) \Theta(\Sigma_{\text{smp}}(\mathbf{r}) - \Sigma_t(\mathbf{r})) \varphi_+(\mathbf{r}, \hat{\Omega}). \quad (\text{A13})$$

In the event that $\Sigma_{\text{smp}}(\mathbf{r}) < \Sigma_t(\mathbf{r})$, negative particles will flip sign, joining the positive particles (V_-). This is modeled in a similar manner, simply flipping the argument in the Heaviside function and replacing the positive angular flux with the negative angular flux:

$$V_- = \Sigma_{\text{smp}}(\mathbf{r}) \Theta(\Sigma_t(\mathbf{r}) - \Sigma_{\text{smp}}(\mathbf{r})) \varphi_-(\mathbf{r}, \hat{\Omega}). \quad (\text{A14})$$

With these terms defined, we may combine our definitions of G_V and $V = V_+ + V_-$ to obtain the gains from virtual collisions, presented in Eq. (A15), having used the function from Eq. (12).

All loss and gain terms have now been defined, leaving us with the final form of the equation, describing the angular flux of the positive particles in Eq. (A16). By

symmetry, the equation for negative particles must have the exact same form and is given in Eq. (A17). Regardless of whether a particle is negative or positive, it is transported in the same manner and by the same rules. These two equations are of course the form presented in Eq. (9) and demonstrate the coupling relationship between positive and negative particles:

$$\begin{aligned}
 G_V &= [1 - q(\mathbf{r})] \frac{|1 - \frac{\Sigma_t(\mathbf{r})}{\Sigma_{\text{smp}}(\mathbf{r})}|}{1 - q(\mathbf{r})} \Sigma_{\text{smp}}(\mathbf{r}) [\Theta(\Sigma_{\text{smp}}(\mathbf{r}) - \Sigma_t(\mathbf{r}))\varphi_+(\mathbf{r}, \hat{\Omega}) + \Theta(\Sigma_t(\mathbf{r}) - \Sigma_{\text{smp}}(\mathbf{r}))\varphi_-(\mathbf{r}, \hat{\Omega})] \\
 &= \Theta(\Sigma_{\text{smp}}(\mathbf{r}) - \Sigma_t(\mathbf{r}))[\Sigma_{\text{smp}}(\mathbf{r}) - \Sigma_t(\mathbf{r})]\varphi_+(\mathbf{r}, \hat{\Omega}) + \Theta(\Sigma_t(\mathbf{r}) - \Sigma_{\text{smp}}(\mathbf{r}))[\Sigma_t(\mathbf{r}) - \Sigma_{\text{smp}}(\mathbf{r})]\varphi_-(\mathbf{r}, \hat{\Omega}) \\
 &= \Delta(\Sigma_{\text{smp}}(\mathbf{r}) - \Sigma_t(\mathbf{r}))\varphi_+(\mathbf{r}, \hat{\Omega}) + \Delta(\Sigma_t(\mathbf{r}) - \Sigma_{\text{smp}}(\mathbf{r}))\varphi_-(\mathbf{r}, \hat{\Omega}), \tag{A15}
 \end{aligned}$$

$$\hat{\Omega} \cdot \nabla \varphi_+ + \Sigma_{\text{smp}} \varphi_+ = \mathcal{S} \varphi_+ + \frac{1}{k} \mathcal{F} \varphi_+ + \Delta(\Sigma_{\text{smp}} - \Sigma_t) \varphi_+ + \Delta(\Sigma_t - \Sigma_{\text{smp}}) \varphi_-, \tag{A16}$$

$$\hat{\Omega} \cdot \nabla \varphi_- + \Sigma_{\text{smp}} \varphi_- = \mathcal{S} \varphi_- + \frac{1}{k} \mathcal{F} \varphi_- + \Delta(\Sigma_{\text{smp}} - \Sigma_t) \varphi_- + \Delta(\Sigma_t - \Sigma_{\text{smp}}) \varphi_+. \tag{A17}$$

-
- [1] H. C. Berg, *Random Walks in Biology* (Princeton University Press, Princeton, 1993).
 - [2] M. S. Bartlett, *Stochastic Population Models in Ecology and Epidemiology* (Methuen, London, 1960).
 - [3] M. M. R. Williams, *Random Processes in Nuclear Reactors*, 1st ed. (Pergamon, New York, 1974).
 - [4] M. F. Modest, *Radiative Heat Transfer* (Academic, New York, 2013).
 - [5] G. I. Bell and S. Glasstone, *Nuclear Reactor Theory* (Van Nostrand Reinhold, New York, 1970).
 - [6] D. M. Arnow, M. H. Kalos, M. A. Lee, and K. E. Schmidt, *J. Chem. Phys.* **77**, 5562 (1982).
 - [7] T. Yamamoto, *Ann. Nucl. Energy* **58**, 72 (2013).
 - [8] T. E. Booth, *Nucl. Sci. Eng.* **143**, 291 (2003).
 - [9] L. Szirmay-Kalos, I. Georgiev, M. Magdics, B. Molnár, and D. Légrády, *Comput. Graph. Forum* **36**, 9 (2017).
 - [10] T. E. Booth and J. E. Gubernatis, *Phys. Rev. E* **80**, 046704 (2009).
 - [11] T. E. Booth and J. E. Gubernatis, *Nucl. Sci. Eng.* **165**, 283 (2010).
 - [12] J. S. Spencer, N. S. Blunt, and W. M. Foulkes, *J. Chem. Phys.* **136**, 054110 (2012).
 - [13] M. H. Holmes, *Introduction to Scientific Computing and Data Analysis*, Texts in Computational Science and Engineering Vol. 13 (Springer, Cham, 2016).
 - [14] I. Lux and L. Koblinger, *Monte Carlo Particle Transport Methods: Neutron and Photon Calculations* (CRC, Boca Raton, 1991).
 - [15] M. Nowak, J. Miao, E. Dumonteil, B. Forget, A. Onillon, K. S. Smith, and A. Zoia, *Ann. Nucl. Energy* **94**, 856 (2016).
 - [16] A. Zoia, E. Dumonteil, and A. Mazzolo, *Phys. Rev. E* **83**, 041137 (2011).
 - [17] C. Demazière, *Int. J. Nucl. Energy Sci. Technol.* **7**, 288 (2013).
 - [18] H. Belanger, D. Mancusi, and A. Zoia, *Eur. Phys. J. Plus* **135**, 877 (2020).
 - [19] C. Demazière, *Modelling of Nuclear Reactor Multi-Physics: From Local Balance Equations to Macroscopic Models in Neutronics and Thermal-Hydraulics* (Academic, New York, 2019).
 - [20] E. Woodcock, T. Murphy, P. Hemmings, and S. Longworth, Techniques used in the GEM code for Monte Carlo neutronics calculations in reactors and other systems of complex geometry, Argonne National Laboratory Report No. ANL-7050, 1965 (unpublished).
 - [21] J. Leppänen, *Ann. Nucl. Energy* **37**, 715 (2010).
 - [22] D. Legrady, B. Molnar, M. Klausz, and T. Major, *Ann. Nucl. Energy* **102**, 116 (2017).
 - [23] B. Molnar, G. Tolnai, and D. Legrady, *Nucl. Sci. Eng.* **190**, 56 (2018).
 - [24] L. L. Carter, E. D. Cashwell, and W. M. Taylor, *Nucl. Sci. Eng.* **48**, 403 (1972).
 - [25] C. Jacoboni and L. Reggiani, *Rev. Mod. Phys.* **55**, 645 (1983).
 - [26] V. Lemaire, M. Thieullen, and N. Thomas, *J. Sci. Comput.* **75**, 1776 (2018).
 - [27] M. Galtier, S. Blanco, C. Caliot, C. Coustet, J. Dauchet, M. El Hafi, V. Eymet, R. Fournier, J. Gautrais, A. Khuong, B. Piaud, and G. Terrée, *J. Quant. Spectrosc. Radiat. Transfer* **125**, 57 (2013).
 - [28] X-5 Monte Carlo Team, MCNP—A general Monte Carlo N-particle transport code, Version 5, Report No. LA-UR-03-1987 (Los Alamos National Laboratory, Los Alamos, 2003), Vol. I.
 - [29] P. K. Romano, N. E. Horelik, B. R. Herman, A. G. Nelson, B. Forget, and K. Smith, *Ann. Nucl. Energy* **82**, 90 (2015).
 - [30] T. E. Booth, A weight (charge) conserving importance-weighted comb for Monte Carlo, Los Alamos National Laboratory Report No. LA-UR 96-0051, 1996 (unpublished).
 - [31] A. Rouchon, A. Zoia, and R. Sanchez, *Ann. Nucl. Energy* **102**, 465 (2017).
 - [32] J. Demmel, in *Proceedings of the IEEE Control Systems Society Workshop on Computer-Aided Control System Design, Tampa, 1989* (IEEE, Piscataway, 1989), pp. 1–7.

- [33] D. A. Brown, M. B. Chadwick, R. Capote, A. C. Kahler, A. Trkov, M. W. Herman, A. A. Sonzogni, Y. Danon, A. D. Carlson, M. Dunn *et al.*, [Nucl. Data Sheets](#) **148**, 1 (2018).
- [34] A. J. M. Plompen, O. Cabellos, C. D. S. Jean, M. Fleming, A. Algora, M. Angelone, P. Archier, E. Bauge, O. Bersillon, A. Blokhin *et al.*, [Eur. Phys. J. A](#) **56**, 181 (2020).
- [35] E. E. Lewis, M. A. Smith, N. Tsoulfanidis, G. Palmiotti, T. A. Taiwo, and R. N. Blomquist, Benchmark specification for Deterministic 2-D/3-D MOX fuel assembly transport calculations without spatial homogenisation (C5G7 MOX), NEA/NSC Report No. JT00105087, 2001 (unpublished).
- [36] C. R. Harris, K. J. Millman, S. J. van der Walt, R. Gommers, P. Virtanen, D. Cournapeau, E. Wieser, J. Taylor, S. Berg, N. J. Smith *et al.*, [Nature \(London\)](#) **585**, 357 (2020).
- [37] H. Belanger, <https://github.com/HunterBelanger/mgmc> (2021).
- [38] J. Rice, *Mathematical Statistics and Data Analysis* (Duxbury, Belmont, 1995).
- [39] T. Yamamoto, [Ann. Nucl. Energy](#) **47**, 14 (2012).