



Original Article

Analysis of inconsistent source sampling in monte carlo weight-window variance reduction methods

David P. Griesheimer*, Virinder S. Sandhu

Naval Nuclear Laboratory, P.O. Box 79, West Mifflin, PA 15122-0079, USA

ARTICLE INFO

Article history:

Received 31 May 2017

Accepted 27 July 2017

Available online 14 August 2017

Keywords:

Monte Carlo

Variance Reduction

Importance Sampling

CADIS

Weight Windows

ABSTRACT

The application of Monte Carlo (MC) to large-scale fixed-source problems has recently become possible with new hybrid methods that automate generation of parameters for variance reduction techniques. Two common variance reduction techniques, weight windows and source biasing, have been automated and popularized by the consistent adjoint-driven importance sampling (CADIS) method. This method uses the adjoint solution from an inexpensive deterministic calculation to define a consistent set of weight windows and source particles for a subsequent MC calculation. One of the motivations for source consistency is to avoid the splitting or rouletting of particles at birth, which requires computational resources. However, it is not always possible or desirable to implement such consistency, which results in inconsistent source biasing. This paper develops an original framework that mathematically expresses the coupling of the weight window and source biasing techniques, allowing the authors to explore the impact of inconsistent source sampling on the variance of MC results. A numerical experiment supports this new framework and suggests that certain classes of problems may be relatively insensitive to inconsistent source sampling schemes with moderate levels of splitting and rouletting.

© 2017 Korean Nuclear Society, Published by Elsevier Korea LLC. This is an open access article under the CC BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/>).

1. Introduction

Many common Monte Carlo (MC) variance reduction techniques rely on weight windows to control the statistical weight of particles during transport in order to minimize the variance of flux or reaction rate tallies in a specified region of phase space (e.g., position, energy, direction). For these techniques, it is well known that the optimal particle weight at each phase location in a problem is given by the objective-driven adjoint flux for that location [1]. Particles with weight outside of a predefined window about the optimal weight are subjected to rouletting or splitting (which adjust the weight in a fair manner) in order to maintain the weight within the weight window. This weight adjustment applies at particle events where the weight changes (e.g., births, collisions) as well as when particles move between regions of phase space with different weight-window parameters.

In early implementations of weight-window variance reduction methods, inconsistencies between the radiation source distribution and the weight window parameters for the simulation resulted in

source particles produced with weights that lay outside of the weight window for the corresponding birth state of the particle. Source particles produced with an inconsistent birth weight are immediately subjected to weight adjustment (splitting or roulette), which is widely believed to decrease the overall effectiveness of the weight-window variance reduction scheme.

In 1998, Wagner and Haight [2] introduced the consistent adjoint-driven importance sampling (CADIS) method for creating adjoint-based sets of weight-window parameters based off of a deterministic estimate for the adjoint flux. In addition, Wagner and Haight showed that the deterministic estimate of the adjoint flux can also be used to define a biased source definition that is consistent with the weight-window parameters. Here, consistent means that source particles from the biased source are born with a weight that lies at the center point of the weight window corresponding to the initial (birth) phase state of the particle. The development of a method for simultaneously creating a consistent source along with the weight-window parameters was a significant advancement and is a major advantage of the CADIS method. The same consistency is also found in the forward-weighted CADIS (FW-CADIS) method used to calculate global MC solutions [3]. However, even with the advancement of the CADIS method, there are some situations where it can be difficult to ensure a completely

* Corresponding author.

E-mail address: david.griesheimer@unnpp.gov (D.P. Griesheimer).

consistent source distribution, and, therefore, the CADIS method cannot be applied as intended.

For example, the biased source produced by CADIS is based on an estimate of the adjoint flux distribution produced from a deterministic solution method – typically the discrete ordinates (S_N) method. As a result, the adjoint flux and the resulting biased source distribution are discretized over space, energy, and direction. In order to reduce the amount of time required to generate weight-window parameters, a relatively coarse discretization may be used to estimate the adjoint flux [3]. Although the discretized biased source produced from CADIS is guaranteed to be consistent with the corresponding weight-window parameters, the CADIS source does not preserve higher-order information about the source distribution, which causes discretization error within the sampled MC source. In practice, many MC codes that use CADIS simply assume that source particles are uniformly distributed within each discretized “bin” of the CADIS source. However, this assumption may still lead to a bias in the results from the MC transport simulation, especially for cases where there is detailed structure in the true source distribution, such as the energy spectrum of a decay source. Any modification of the source distribution to reduce this bias may lead to inconsistencies between the adjusted source and the original weight-window parameters.

In other situations, the radiation source may be “presampled” from a preceding calculation and stored as a census file containing detailed state information about the source particles. This scenario is common when generating secondary radiations during MC transport (e.g., (n,γ) or (γ,n) reactions), exchanging information between MC eigenvalue and fixed-source calculations, or in S_N /MC or MC/MC splice calculations where particles that reach a pre-defined “trapping surface” are stored for use in a subsequent MC simulation [4]. In these cases, it is straightforward to collapse the particle census into a discretized source representation for use in CADIS. However, replacing the census source by the discretized representation would eliminate valuable information stored in the census, such as the correlations between the position, energy, and direction of each particle, and is not a practical solution for many splice calculations. Therefore, retaining the particle census introduces inconsistent source sampling into the subsequent MC calculation.

In addition, for some types of analyses, it is desirable to generate a single set of weight-window parameters that can be used with a range of similar model configurations, often representing source, geometry, or composition perturbations with respect to a single reference scenario. In these cases, the CADIS method is well suited for determining the weight-window parameters and a consistent source for the reference configuration, but it can become expensive if the weight-window parameters and/or the consistent biased source must be regenerated for every model perturbation. In practice, a single set of weight-window parameters is often used for all of the model perturbations, regardless of whether each source distribution is actually consistent with the weight-window parameters.

Finally, we note that, although the CADIS method has proven to be extremely successful, there are still weight-window variance reduction techniques in use [5,6], and under development, that do not produce a consistent biased source distribution, for a variety of reasons.

For any situation where an inconsistent source distribution may be used with weight-window variance reduction, it is important to have a clear understanding of the effects of weight adjustment via splitting or rouletting immediately after particle birth. Although the conventional wisdom maintains that any weight adjustment at birth will reduce the effectiveness of the weight-window variance reduction, no systematic, formal investigation of this conjecture has

ever been performed to our knowledge. Although it appears self-evident that frequently adjusting the initial weight of source particles is counterproductive, it seems reasonable that a small initial weight adjustment for source particles may be acceptable for many applications. However, such a conclusion requires a thorough characterization of the effect of inconsistent source sampling based on the degree of inconsistency.

In this paper, we develop a mathematical framework for quantifying the impact of inconsistent source sampling on the variance of tallied quantities in a MC simulation. The derived relationships are supported with results from numerical experiments and provide a foundation for additional analyses tailored to a variety of specific applications.

2. Expected variance by sample scheme

In this section, we derive the expected variance in estimated response for several different source sampling schemes. Prior to proceeding, it is useful to define notation and significant statistical relationships that will be used throughout the remainder of the paper.

2.1. Notation and basic relationships

In MC transport methods, each history can be viewed as the combination of two separate realizations: the initial (birth) state of the source particle, denoted \mathbf{x} , and the response of the history as measured against some predetermined objective, denoted r . In this context, we have assumed that the initial particle state, \mathbf{x} , is a vector that includes properties such as the birth energy, position, and direction of the particle, and that the response, r , is a scalar value. Note that these are arbitrary assumptions and may be changed without loss of generality.

To an external observer, ignorant of the inner workings of the MC transport algorithm, it appears that each history produces a realization (\mathbf{x}, r) from the joint probability density function (PDF) $p(\mathbf{x}, r)$. Based on the properties of joint probability distributions, it follows that the expected value and variance of any function $f(\mathbf{x}, r)$ applied to a realization of the joint PDF is given by

$$E[f(\mathbf{x}, r)] = \int_{-\infty}^{\infty} \int_{\Gamma} f(\mathbf{x}, r) p(\mathbf{x}, r) d\mathbf{x} dr \quad (1)$$

and

$$\text{Var}[f(\mathbf{x}, r)] = E[f^2(\mathbf{x}, r)] - E[f(\mathbf{x}, r)]^2, \quad (2)$$

where Γ is the domain for the birth state of the particle.

Note that the joint PDF $p(\mathbf{x}, r)$ can be written as the product of conditional and marginal probability distributions, $p(\mathbf{x}, r) = p(r|\mathbf{x})p(\mathbf{x})$. In this case, the expected value of the function $f(\mathbf{x}, r)$ can be expressed as

$$E[f(\mathbf{x}, r)] = E_{\mathbf{x}}[E_r[f(\mathbf{x}, r)|\mathbf{x}]] = \int_{\Gamma} E_r[f(\mathbf{x}, r)|\mathbf{x}] p(\mathbf{x}) d\mathbf{x}, \quad (3)$$

where

$$E_r[f(\mathbf{x}, r)|\mathbf{x}] = \int_{-\infty}^{\infty} f(\mathbf{x}, r) p(r|\mathbf{x}) dr, \quad (4)$$

and subscripts have been included on the expectation operators to clarify which variable the expectation is taken with respect to. The

relationship in Eq. (3) is commonly referred to as the fundamental property of conditional expected values, the law of total expectation, or the law of iterated expectation.

Similarly, the law of total variance can be used to express the variance of $f(\mathbf{x}, r)$ in terms of the conditional probability,

$$\text{Var}[f(\mathbf{x}, r)] = E_{\mathbf{x}}[\text{Var}[f(\mathbf{x}, r)|\mathbf{x}]] + \text{Var}[E_r[f(\mathbf{x}, r)|\mathbf{x}]], \quad (5)$$

where

$$\text{Var}[f(\mathbf{x}, r)|\mathbf{x}] = E_r \left[(f(\mathbf{x}, r) - E_r[f(\mathbf{x}, r)|\mathbf{x}])^2 | \mathbf{x} \right]. \quad (6)$$

Note that Eq. (5) demonstrates that the total variance of $f(\mathbf{x}, r)$ includes two components: the first term gives the variance due only to the randomness of r for a fixed value of \mathbf{x} , referred to as the transport variance, and the second term gives the variance due to the effect of the randomness of \mathbf{x} on the conditional distribution $p(r|\mathbf{x})$, which is referred to as the source variance. The use of the conditional probability is often used for analysis of MC radiation transport algorithms because it explicitly shows the dependence of the response on the initial state of a particle.

Now that the notation and relevant statistical relationships have been established, we will consider the variance of the estimated response for several different source sampling schemes.

2.2. Unbiased sampling scheme

Consider a fair MC transport process where each particle history i begins with an initial particle state $\tilde{\mathbf{x}}_i$ sampled from the distribution $p(\mathbf{x})$ and produces a corresponding response \tilde{r}_i according to the probability distribution $p(r|\mathbf{x})$. As described in the previous section, the initial state and corresponding response can be viewed as either sequential (conditional) realizations, or as a single realization of the ordered pair $(\tilde{\mathbf{x}}_i, \tilde{r}_i)$ from the joint PDF $p(\mathbf{x}, r)$.

For a simulation with N independent particle histories, an estimator for the expected response can be computed with the unbiased sample statistic

$$\hat{\mu}_r = \frac{1}{N} \sum_{i=1}^N \tilde{r}_i. \quad (7)$$

The variance of the response can be estimated using the sample variance statistic

$$\hat{s}_r^2 = \frac{1}{N-1} \sum_{i=1}^N (\tilde{r}_i - \hat{\mu}_r)^2, \quad (8)$$

which is an unbiased estimator for the variance of the response r ,

$$\text{Var}[r] = E_{\mathbf{x}}[\text{Var}[r|\mathbf{x}]] + \text{Var}[E_r[r|\mathbf{x}]] \quad (9)$$

by the law of total variance [Eq. (5)].

In this context, the first term in Eq. (9) has the physical interpretation as the variance due to the randomness of the transport process (transport variance), whereas the second term is the variance due to the randomness of the initial birth state of each source particle (source variance).

Finally, we note that the variance of the sum of uncorrelated variables is equal to the sum of the variance of the individual random variables (a relationship referred to in some texts as the Bienaymé formula), which allows the variance of the estimator for the mean response $\hat{\mu}_r$ to be written as

$$\text{Var}[\hat{\mu}_r] = \text{Var} \left[\frac{1}{N} \sum_{i=1}^N r_i \right] = \frac{1}{N^2} \text{Var} \left[\sum_{i=1}^N r_i \right] = \frac{\text{Var}[r]}{N}. \quad (10)$$

Substituting Eq. (9) into Eq. (10) yields the final expression

$$\text{Var}[\hat{\mu}_r] = \frac{E_{\mathbf{x}}[\text{Var}[r|\mathbf{x}]]}{N} + \frac{\text{Var}[E_r[r|\mathbf{x}]]}{N}. \quad (11)$$

2.3. Importance sampling scheme

The previous section established the basic estimators for the expected response and associated variance using only unbiased realizations from the joint distribution $p(\mathbf{x}, r)$. However, it is well known that the variance of the expected response can be reduced via importance sampling [7]. In adjoint-driven importance sampling, the underlying source distribution (in this case, $p(\mathbf{x})$) is altered so that the initial particle state is drawn from a distribution that is proportional to the corresponding response of the source. In order to preserve the original source distribution, the observed response of the particles is weighted in proportion to the ratio of the probabilities of the initial state in the original and modified probability distributions.

For the MC radiation transport process, let us define a biased source distribution $p'(\mathbf{x})$ such that

$$p'(\mathbf{x}, r) = p(r|\mathbf{x})p'(\mathbf{x}). \quad (12)$$

The corresponding weighting factor for each state point is given by

$$w(\mathbf{x}) = p(\mathbf{x})/p'(\mathbf{x}). \quad (13)$$

With importance sampling, the unbiased statistic for estimating the expected response is

$$\hat{\mu}_r^{\text{imp}} = \frac{1}{N} \sum_{i=1}^N w(\tilde{\mathbf{x}}'_i) \tilde{r}_i, \quad (14)$$

where $\tilde{\mathbf{x}}'_i$ denotes a source state realization sampled from the modified distribution $p'(\mathbf{x})$.

Similarly, the sample variance statistic

$$\hat{s}_{r,\text{imp}}^2 = \frac{1}{N-1} \sum_{i=1}^N (w(\tilde{\mathbf{x}}'_i) \tilde{r}_i - \hat{\mu}_r^{\text{imp}})^2, \quad (15)$$

is an unbiased estimator for $\text{Var}[w(\mathbf{x}')r]$, where

$$\text{Var}[w(\mathbf{x}')r] = E_{\mathbf{x}'}[\text{Var}[w(\mathbf{x}')r|\mathbf{x}']] + \text{Var}[E_r[w(\mathbf{x}')r|\mathbf{x}']]. \quad (16)$$

Recognizing that the weight factor $w(\mathbf{x}')$ is a constant with respect to the variance of $r|\mathbf{x}'$, the first term on the right-hand side of Eq. (16) can be rewritten as

$$E_{\mathbf{x}'}[\text{Var}[w(\mathbf{x}')r|\mathbf{x}']] = E_{\mathbf{x}'}[w^2(\mathbf{x}')\text{Var}[r|\mathbf{x}']]. \quad (17)$$

Applying the definition for the expectation operator $E_{\mathbf{x}'}$ to Eq. (17) yields

$$E_{\mathbf{x}'}[\text{Var}[w(\mathbf{x}')r|\mathbf{x}']] = \int_{\Gamma} \text{Var}[r|\mathbf{x}'] w^2(\mathbf{x}') p'(\mathbf{x}') d\mathbf{x}'. \quad (18)$$

Using Eq. (13) to relate $p(\mathbf{x}') = w(\mathbf{x}')p'(\mathbf{x}')$ and switching the dummy variable of integration from \mathbf{x}' to \mathbf{x} allows Eq. (18) to be rewritten as

$$E_{\mathbf{x}'}[\text{Var}[w(\mathbf{x}')r|\mathbf{x}']] = \int_{\Gamma} w(\mathbf{x})\text{Var}[r|\mathbf{x}]p(\mathbf{x})d\mathbf{x}, \quad (19)$$

which can be written in expectation notation as

$$E_{\mathbf{x}'}[\text{Var}[w(\mathbf{x}')r|\mathbf{x}']] = E_{\mathbf{x}}[w(\mathbf{x})\text{Var}[r|\mathbf{x}]]. \quad (20)$$

Substituting Eq. (18) into Eq. (16) gives an expression for the response variance under importance sampling.

$$\text{Var}[w(\mathbf{x}')r] = E_{\mathbf{x}}[w(\mathbf{x})\text{Var}[r|\mathbf{x}]] + \text{Var}[w(\mathbf{x}')E_r[r|\mathbf{x}']]. \quad (21)$$

Note that the first term on the right-hand side of Eq. (21) is expressed in terms of the expectation and variance with respect to the original source distribution, $p(\mathbf{x})$, rather than the modified distribution $p'(\mathbf{x})$.

Finally, applying the Bienaymé formula for the variance of the sum of uncorrelated variables to Eq. (14) and using the result for $\text{Var}[w(\mathbf{x}')r]$ gives the final variance for the statistic $\hat{\mu}_r$

$$\text{Var}[\hat{\mu}_r^{\text{imp}}] = \frac{E_{\mathbf{x}}[w(\mathbf{x})\text{Var}[r|\mathbf{x}]]}{N} + \frac{\text{Var}[w(\mathbf{x}')_r[r|\mathbf{x}']]}{N}. \quad (22)$$

Comparing Eqs. (22) and (11) illustrates the effect of importance sampling via the presence of the weight parameter $w(\mathbf{x})$ in each term. It is particularly interesting to note that the weighting parameter affects both the transport variance and source variance terms. The effect on the transport variance term is somewhat surprising, as it suggests that importance sampling schemes that emphasize initial birth states with above-average response (transport) variance may diminish the effectiveness of importance sampling for reducing total variance.

The objective of the importance sampling scheme described in this section is to reduce the total variance of the response by altering the distribution of source states in an unbiased way. When importance sampling is applied to a process that involves a probability distribution of only one variable, say $p(\mathbf{x})$, it is known that the optimal weighting function for calculating the expected value of any response function $g(\mathbf{x})$ is given by $w(\mathbf{x}) = E[g(\mathbf{x})]/g(\mathbf{x})$. In fact, this optimal weight gives $\text{Var}[w(\mathbf{x})g(\mathbf{x})] = 0$ by ensuring that each sample gives a response that is exactly equal to the expected response $E[g(\mathbf{x})]$.

However, the identification of an optimal weighting function for importance sampling is not as straightforward when dealing with a joint probability distribution. By analogy with the single-variable case, it seems reasonable to assume that the weight should be defined as

$$w(\mathbf{x}) = \frac{\int_{\Gamma} E_r[r|\mathbf{x}]p(\mathbf{x})d\mathbf{x}}{E_r[r|\mathbf{x}]}, \quad (23)$$

which is the ratio of the mean total response to the mean response conditioned on the initial state \mathbf{x} . In fact, it can be immediately shown that the weighting function defined in Eq. (23) optimizes the variance associated with the selection of an initial source state by eliminating the source variance term ($\text{Var}[E_r[w(\mathbf{x}')r|\mathbf{x}']] = 0$). Again, this is because the weighting function ensures that each source state contributes exactly the same expected response.

Applying the weight function in Eq. (23) to the equation for total variance of the response, Eq. (21), yields

$$\frac{\text{Var}[w(\mathbf{x})r]}{E[r]} = \int_{\Gamma} \left(\frac{\text{Var}[r|\mathbf{x}]}{E_r[r|\mathbf{x}]} \right) p(\mathbf{x})d\mathbf{x}, \quad (24)$$

which shows that the relative variance of the total response is equal to the average of the relative variance for all of the conditional responses taken with respect to the true source distribution $p(\mathbf{x})$.

Note that the use of importance sampling for the source state does not eliminate the variability in observed response owing to the transport process (transport variance). In order to achieve a true zero-variance process, it is necessary to also adjust the transport process so that each initial state has the same observed response, not just the same expected response as shown above.

In consistent-source variance reduction methods, the source weighting function is chosen based on Eq. (23), in order to minimize (or eliminate) the source variance term. It follows that this importance weighting function does not necessarily minimize the total variance of the response, as demonstrated by the example problem described in the results section. For simplified problems, it is possible to show that a weighting function obtained by minimizing both the transport and source variance terms in Eq. (21) produces a lower total variance. Generalized conditions for the optimum weighting function have not been derived. However, it should be noted that results from preliminary testing covering a range of possible model conditions suggest that the zero-source-variance weighting function given in Eq. (23) produces responses within a few percent of the optimal variance for realistic situations.

2.4. Source splitting scheme

Next, let us consider a generalization of the expected response estimator in which source particles are split at birth and multiple independent realizations are generated for each initial state point.

For an MC simulation with N' independently sampled state points and M independent realizations for each initial state point, the unbiased importance sampling statistic for estimating the expected response is

$$\hat{\mu}_r^{\text{split}} = \frac{1}{N'} \sum_{i=1}^{N'} \frac{1}{M} \sum_{j=1}^M w(\mathbf{x}'_i) \bar{r}_{j,i} \equiv \frac{1}{N'} \sum_{i=1}^{N'} \zeta_{M,i}^{\text{split}}, \quad (25)$$

where $\bar{r}_{j,i}$ is a realization from the conditional distribution $p(r|\mathbf{x}'_i)$, and $\zeta_{M,i}^{\text{split}}$ is a realization of the random variable obtained by taking the sample mean of the responses due to transporting M replicates of initial source state \mathbf{x}'_i . Note that the variable $\zeta_{M,i}^{\text{split}}$ is actually a function of a realization taken from the joint probability distribution $p(\mathbf{x}, r_1, \dots, r_M)$. Such a realization is not physically realistic, because it implies that a single source site will induce M response realizations. However, the joint realization is an accurate representation of MC particle transport with source particle splitting. Furthermore, we note that the response variables r_1, \dots, r_M are conditionally independent, which means that each variable is conditionally dependent on the random variable \mathbf{x} , but independent from all of the other response variables. The property of conditional independence allows the joint PDF for the birth state and M responses to be written as

$$p(\mathbf{x}, r_1, \dots, r_M) = p(r_1, \dots, r_M|\mathbf{x})p(\mathbf{x}) = p(\mathbf{x}) \prod_{j=1}^M p(r_j|\mathbf{x}). \quad (26)$$

Further recognizing that all of the response realizations are taken from a common conditional PDF, $p(r_j|\mathbf{x}) = p(r|\mathbf{x})$ for all r_j , allows the joint PDF in Eq. (26) to be written as

$$p(\mathbf{x}, r_1, \dots, r_M) = p(\mathbf{x})(p(r|\mathbf{x}))^M. \quad (27)$$

Returning to the expected response estimator, $\hat{\mu}_r^{\text{split}}$, defined in Eq. (25), it follows that the corresponding sample variance statistic

$$\hat{S}_{r,\text{split}}^2 = \frac{1}{N' - 1} \sum_{i=1}^{N'} \left(\frac{1}{M} \sum_{j=1}^M w(\mathbf{x}'_i) \tilde{r}_{j,i} - \hat{\mu}_r^{\text{split}} \right)^2, \quad (28)$$

is an unbiased estimator for $\text{Var}[\zeta_M^{\text{split}}]$, where

$$\text{Var}[\zeta_M^{\text{split}}] = E_{\mathbf{x}'}[\text{Var}[\zeta_M^{\text{split}}|\mathbf{x}']] + \text{Var}[E_r[\zeta_M^{\text{split}}|\mathbf{x}']]. \quad (29)$$

Fortunately, Eq. (29) can be recast into a more meaningful form. We begin by expanding the first term (transport variance) on the right-hand side of Eq. (29)

$$E_{\mathbf{x}'}[\text{Var}[\zeta_M^{\text{split}}|\mathbf{x}']] = E_{\mathbf{x}'}\left[\text{Var}\left[\frac{1}{M} \sum_{j=1}^M w(\mathbf{x}') r_j \middle| \mathbf{x}'\right]\right]. \quad (30)$$

Recognizing that the number of replicates, M , and the weight factor $w(\mathbf{x})$ are independent of the replicate number and response variance allows Eq. (30) to be rewritten as

$$E_{\mathbf{x}'}[\text{Var}[\zeta_M^{\text{split}}|\mathbf{x}']] = E_{\mathbf{x}'}\left[\frac{w^2(\mathbf{x}')}{M^2} \text{Var}\left[\sum_{j=1}^M r_j \middle| \mathbf{x}'\right]\right]. \quad (31)$$

Because the realizations of r_{ij} (i.e., $r_j|\mathbf{x}$) are conditionally independent, it is possible to apply the Bienaymé formula to write

$$E_{\mathbf{x}'}[\text{Var}[\zeta_M^{\text{split}}|\mathbf{x}']] = \frac{E_{\mathbf{x}'}[w^2(\mathbf{x}') \text{Var}[r|\mathbf{x}']]}{M}. \quad (32)$$

Finally, we apply the procedure outlined in Eqs. (18–20) to write Eq. (32) in terms of the expected value with respect to the unbiased source distribution, $p(\mathbf{x})$

$$E_{\mathbf{x}'}[\text{Var}[\zeta_M^{\text{split}}|\mathbf{x}']] = \frac{E_{\mathbf{x}}[w(\mathbf{x}) \text{Var}[r|\mathbf{x}]]}{M}. \quad (33)$$

Returning to the second term (source variance) in Eq. (29), we can again expand this term and factor out constants to give

$$\text{Var}[E_r[\zeta_M^{\text{split}}|\mathbf{x}']] = \text{Var}\left[\frac{w(\mathbf{x}')}{M} E_r\left[\sum_{j=1}^M (r_j|\mathbf{x}') \middle| \mathbf{x}'\right]\right]. \quad (34)$$

Applying the definition of the conditional expected value for the response, Eq. (4), gives

$$\text{Var}[E_r[\zeta_M^{\text{split}}|\mathbf{x}']] = \text{Var}\left[\frac{w(\mathbf{x}')}{M} \sum_{j=1}^M \int_0^\infty r_j p(r_j|\mathbf{x}') dr_j\right]. \quad (35)$$

Because the random variables r_j are conditionally independent and drawn from the same probability distribution $p(r|\mathbf{x})$, it follows that

$$\begin{aligned} \text{Var}[E_r[\zeta_M^{\text{split}}|\mathbf{x}']] &= \text{Var}\left[w(\mathbf{x}') \mathbf{x}' \int_0^\infty r p(r|\mathbf{x}') dr\right] \\ &= \text{Var}[w(\mathbf{x}') E_r[r|\mathbf{x}']]. \end{aligned} \quad (36)$$

Substituting Eqs. (33) and (36) into Eq. (29), gives the final simplified expression for the variance of ζ_M^{split}

$$\text{Var}[\zeta_M^{\text{split}}] = \frac{E_{\mathbf{x}}[w(\mathbf{x}) \text{Var}[r|\mathbf{x}]]}{M} + \text{Var}[w(\mathbf{x}') E_r[r|\mathbf{x}']]. \quad (37)$$

Inspection of Eq. (37) indicates that source splitting only affects the transport variance term, and that the total response variance per source sample will decrease as the splitting factor M increases.

Equation (37) gives the variance for the mean response of M replicates for a single source state. The variance for the mean response over N' independent trials, $\hat{\mu}_r^{\text{split}}$, can be determined by applying the Bienaymé formula to Eq. (25), then substituting the expression in Eq. (37), to yield

$$\text{Var}[\hat{\mu}_r^{\text{split}}] = \frac{E_{\mathbf{x}}[\text{Var}[w(\mathbf{x}) r|\mathbf{x}]]}{MN'} + \frac{\text{Var}[w(\mathbf{x}') E_r[r|\mathbf{x}']]}{N'}. \quad (38)$$

A comparison of Eq. (38) with the variance for importance sampling without splitting [Eq. (22)] shows that the only difference between the expressions is the presence of the $1/M$ factor in the first term of Eq. (38). Note that if the number of source samples are held constant between importance sampling with and without source splitting (e.g., $N' = N$), source splitting will cause the overall response variance to go down, because of a decrease in the transport variance term. This makes sense, as the source replicates are reducing the statistical uncertainty in the response associated with transporting radiation from each sampled initial particle state.

However, if N' and M are constrained such that the total amount of transport work is held constant ($MN' = N$), a question arises regarding optimal allocation of resources between M and N' . In cases where the source variance is larger than the transport variance, the optimal allocation is to maximize the number of source particles, $N' = N$, and use only one replicate for each source. However, as the magnitude of the source variance becomes small relative to the transport variance, the penalty (i.e., increase in total variance) associated with source splitting decreases. In the limiting case, when the source variance is equal to zero (i.e., all weighted source sites produce the same expected response), the total response variance is not affected by source splitting.

This observation leads to several interesting conclusions. First, this result confirms the long-held conventional wisdom that source splitting will increase the variance of tallied quantities in simulations for which it is used, assuming that total work is held constant. However, the result also shows that the increase in variance due to splitting may be small in cases where the source variance is small relative to the transport variance. This second result is especially intriguing when recalling that importance sampling is used to minimize the source variance term. Thus, in this regime, source splitting may be used to correct for an inconsistent source distribution without a significant increase in total response variance.

2.5. Source rouletting scheme

Let us now consider an expected response estimator in which source particles are subjected to Russian roulette at birth. For an MC simulation with N' independently sampled state points, the unbiased importance sampling statistic for estimating the expected response with Russian roulette is

$$\hat{\mu}_r^{\text{roulette}} = \frac{1}{N'} \sum_{i=1}^{N'} \frac{\tilde{b}_i w(\tilde{\mathbf{x}}'_i) \tilde{r}_i}{M} \equiv \frac{1}{N'} \sum_{i=1}^{N'} \zeta_{M,i}^{\text{roulette}}, \quad (39)$$

where \tilde{r}_i and $\tilde{\mathbf{x}}'_i$ are realizations from the joint distribution $p(\mathbf{x}', r)$, and \tilde{b}_i is a realization from a Bernoulli distribution with probability of success given by M (for $M \leq 1$), and $\zeta_{M,i}^{\text{roulette}}$ is a realization of the random variable obtained by setting the response equal to zero with probability $(1 - M)$. Note that the random variable b

determines whether each particle will survive the roulette process, and is completely independent from the random variables \mathbf{x} and r .

In this case, the sample variance statistic

$$\hat{S}_{r,\text{roulette}}^2 = \frac{1}{N' - 1} \sum_{i=1}^{N'} \left(\frac{\tilde{b}_i w(\tilde{\mathbf{x}}'_i) \tilde{r}_i}{M} - \hat{\mu}_r^{\text{roulette}} \right)^2, \quad (40)$$

is an unbiased estimator for $\text{Var}[\zeta_M^{\text{roulette}}]$, where

$$\text{Var}[\zeta_M^{\text{roulette}}] = E \left[\left(\frac{bw(\mathbf{x}')r}{M} \right)^2 \right] - E \left[\frac{bw(\mathbf{x}')r}{M} \right]^2, \quad (41)$$

which can be factored to yield

$$\text{Var}[\zeta_M^{\text{roulette}}] = \frac{1}{M^2} \left(E[b^2] E[(w(\mathbf{x}')r)^2] - E[b]^2 E[w(\mathbf{x}')r]^2 \right), \quad (42)$$

because the realizations of (\mathbf{x}', r) are independent from the realizations of b . Recognizing that $E[b^2] = M$ and $E[b]^2 = M^2$ for a Bernoulli distributed random variable enables Eq. (42) to be simplified, giving

$$\text{Var}[\zeta_M^{\text{roulette}}] = \frac{1}{M} \text{Var}[w(\mathbf{x}')r] + \frac{1-M}{M} E[r]^2. \quad (43)$$

Again, the variance for the mean response over N' independent trials, $\hat{\mu}_r^{\text{roulette}}$, can be determined by applying the Bienaymé formula to Eq. (39), then substituting the expression in Eq. (43) to yield

$$\text{Var}[\hat{\mu}_r^{\text{roulette}}] = \frac{1}{MN'} \text{Var}[w(\mathbf{x}')r] + \frac{1-M}{MN'} E[r]^2. \quad (44)$$

For comparison with the corresponding variance for importance sampling with [Eq. (38)] and without splitting [Eq. (38)], the previous result for $\text{Var}[w(\mathbf{x}')r]$ given in Eq. (21) can be applied to Eq. (44), yielding an alternate form

$$\text{Var}[\hat{\mu}_r^{\text{roulette}}] = \frac{E_{\mathbf{x}}[w(\mathbf{x})\text{Var}[r|\mathbf{x}]]}{MN'} + \frac{\text{Var}[w(\mathbf{x}')E_r[r|\mathbf{x}']]}{MN'} + \frac{1-M}{MN'} E[r]^2. \quad (45)$$

From Eq. (44), it is clear that applying Russian roulette will always cause the response variance to increase relative to importance sampling for a fixed number of source samples ($N' = N$). This increase is attributable both to a decrease in the denominator of each term because $M < 1$, as well as the presence of an extra additive term proportional to the square of the expected response.

If N' and M are constrained such that the total amount of transport work is held constant ($M N' = N$), Eq. (44) shows that rouletting will always produce an increase in response variance relative to a traditional importance sampling scheme [Eq. (22)]. This variance penalty appears as an additive term, which is proportional to $(1 - M) \cdot E[r]^2$. Unlike the splitting scheme considered previously, which was dependent on the biased source distribution used for importance sampling, the variance penalty for rouletting only depends on the roulette survival probability and the expected response. Interestingly, Eq. (44) indicates that there is no increase in variance when rouletting particles in a system where the expected response is zero.

2.6. Generalized weight adjustment scheme

In Sections 2.4 and 2.5, formulations for the response variance in the presence of source splitting (for an integer split ratio) and source rouletting were presented, respectively. In this section, we derive a

common response estimator that can be used for both source rouletting and splitting, including noninteger splitting ratios.

For an MC simulation with N' independently sampled state points, the generalized sampling statistic for estimating the expected response is

$$\hat{\mu}_r^{\text{general}} = \frac{1}{N'} \sum_{i=1}^{N'} \frac{1}{M} \sum_{j=1}^{[M]} \tilde{b}'_j w(\tilde{\mathbf{x}}'_i) \tilde{r}_{j,i} \equiv \frac{1}{N'} \sum_{i=1}^{N'} \zeta_{M,i}^{\text{general}}, \quad (46)$$

where $\tilde{r}_{j,i}$ is a realization from the conditional distribution $p(r|\mathbf{x}'_i)$, \tilde{b}'_j is a realization from a Bernoulli distribution with probability of success given by $M/[M]$, $\zeta_{M,i}^{\text{general}}$ is a realization of the random variable obtained by taking the sample mean of the responses due to transporting $[M]$ replicates of the initial source site $\tilde{\mathbf{x}}'_i$, and the symbol $[\cdot]$ denotes a ceiling operation such that $[M]$ is the smallest integer that is larger than M . As described in Section 2.4, the variable $\zeta_{M,i}^{\text{general}}$ is actually a function of a realization taken from the conditionally independent joint probability distribution $p(\mathbf{x}, r_1, \dots, r_M)$. In addition, the Bernoulli realizations \tilde{b}'_j are assumed to be unconditionally independent from the random variables \mathbf{x} and r .

Note that in Eq. (46) the parameter M determines how the weight adjustment is performed, with $M < 1$ causing source rouletting, $M > 1$ causing source splitting, and $M = 1$ reverting to standard importance sampling. It is also important to recognize that the estimator defined in Eq. (46) automatically accounts for noninteger splitting ratios by adjusting the effective splitting ratio upward to the integer value $[M]$ and then applying rouletting to eliminate the excess particles produced during splitting.

Based on the form of Eq. (46), it follows that the corresponding sample variance statistic

$$\hat{S}_{r,\text{general}}^2 = \frac{1}{N' - 1} \sum_{i=1}^{N'} \left(\frac{1}{M} \sum_{j=1}^{[M]} \tilde{b}'_j w(\tilde{\mathbf{x}}'_i) \tilde{r}_{j,i} - \hat{\mu}_r^{\text{general}} \right)^2 \quad (47)$$

is an unbiased estimator for $\text{Var}[\zeta_M^{\text{general}}]$, where

$$\text{Var}[\zeta_M^{\text{general}}] = E_{\mathbf{x}'} \left[\text{Var}[\zeta_M^{\text{general}}|\mathbf{x}'] \right] + \text{Var} \left[E_r[\zeta_M^{\text{general}}|\mathbf{x}'] \right]. \quad (48)$$

Again, Eq. (48) can be simplified into a more intuitive form, following the same general approach as applied in Section 2.4.

As usual, we begin by expanding the first term on the right-hand side of Eq. (48) and moving constant factors out of the variance operator, which gives

$$E_{\mathbf{x}'} \left[\text{Var}[\zeta_M^{\text{general}}|\mathbf{x}'] \right] = E_{\mathbf{x}'} \left[\frac{w^2(\mathbf{x}')}{M^2} \text{Var} \left[\sum_{j=1}^{[M]} \tilde{b}'_j r_j \middle| \mathbf{x}' \right] \right]. \quad (49)$$

Recognizing that the realizations from b' are independent and the r_j variables are conditionally independent allows application of the Bienaymé formula to Eq. (49), resulting in

$$E_{\mathbf{x}'} \left[\text{Var}[\zeta_M^{\text{general}}|\mathbf{x}'] \right] = E_{\mathbf{x}'} \left[\frac{w^2(\mathbf{x}') [M]}{M^2} \text{Var}[b'r|\mathbf{x}'] \right]. \quad (50)$$

By the definition of variance and the independence of the random variables b' and r , it follows that

$$\text{Var}[b'r|\mathbf{x}'] = E[b'^2] E[r^2|\mathbf{x}'] - (E[b'] E[r|\mathbf{x}'])^2. \quad (51)$$

For a Bernoulli distribution with probability of success $M/|M|$, it is straightforward to show that the expected value of b' and b'^2 are $E[b'] = E[b'^2] = M/|M|$. Substituting these values and Eq. (51) into Eq. (50) and simplifying gives

$$E_{\mathbf{x}'} \left[\text{Var} \left[\zeta_M^{\text{general}} | \mathbf{x}' \right] \right] = E_{\mathbf{x}'} \left[\frac{w^2(\mathbf{x}')}{M} \left(E[r^2 | \mathbf{x}'] - \left(\frac{M}{|M|} \right) E[r | \mathbf{x}']^2 \right) \right]. \quad (52)$$

Again, we apply the procedure outlined in Eqs. (18–20) to write Eq. (52) in terms of the expected value with respect to the unbiased source distribution, $p(\mathbf{x})$,

$$E_{\mathbf{x}'} \left[\text{Var} \left[\zeta_M^{\text{general}} | \mathbf{x}' \right] \right] = E_{\mathbf{x}} \left[\frac{w(\mathbf{x})}{M} \left(E[r^2 | \mathbf{x}] - \left(\frac{M}{|M|} \right) E[r | \mathbf{x}]^2 \right) \right]. \quad (53)$$

Returning Eq. (48), expanding the second term, and moving constant factors outside of the expectation operator gives

$$\text{Var} \left[E_r \left[\zeta_M^{\text{general}} | \mathbf{x}' \right] \right] = \text{Var} \left[\frac{w(\mathbf{x}')}{M} \sum_{j=1}^{|M|} E_r [b' r_j | \mathbf{x}'] \right]. \quad (54)$$

Because the random variables b' and r_j are independent and all of the r_j variables are sampled from a common probability distribution, it follows that

$$\text{Var} \left[E_r \left[\zeta_M^{\text{general}} | \mathbf{x}' \right] \right] = \text{Var} \left[\frac{w(\mathbf{x}') |M|}{M} E[b'] E_r[r | \mathbf{x}'] \right]. \quad (55)$$

Since $E[b'] = M/|M|$, Eq. (55) reduces to

$$\text{Var} \left[E_r \left[\zeta_M^{\text{general}} | \mathbf{x}' \right] \right] = \text{Var} [w(\mathbf{x}') E_r[r | \mathbf{x}']]. \quad (56)$$

Substituting Eqs. (53) and (55) into Eq. (48) gives the final simplified expression for the variance of ζ_M^{general}

$$\text{Var} \left[\zeta_M^{\text{general}} \right] = E_{\mathbf{x}} \left[w(\mathbf{x}) \left(\frac{E[r^2 | \mathbf{x}]}{M} - \left(\frac{1}{|M|} \right) E[r | \mathbf{x}]^2 \right) \right] + \text{Var} [w(\mathbf{x}') E_r[r | \mathbf{x}']]. \quad (57)$$

Again, the variance for the mean response over N' independent trials, $\hat{\mu}_r^{\text{general}}$, can be determined by applying the Bienaymé formula to Eq. (46), then substituting the expression in Eq. (57), giving

$$\text{Var} \left[\hat{\mu}_r^{\text{general}} \right] = E_{\mathbf{x}} \left[w(\mathbf{x}) \left(\frac{E[r^2 | \mathbf{x}]}{MN'} - \left(\frac{1}{|M|N'} \right) E[r | \mathbf{x}]^2 \right) \right] + \frac{\text{Var} [w(\mathbf{x}') E_r[r | \mathbf{x}']]}{N'}. \quad (58)$$

It is possible to show that the variance for the generalized weight adjustment estimator [Eq. (58)] reduces to: the variance of the source splitting estimator [Eq. (38)] when M is an integer value > 1 , the variance of the roulette estimator [Eq. (44)] when $M < 1$, and the variance of the original importance sampling estimator [Eq. (22)] when $M = 1$, as expected.

3. Inconsistent source sampling

In order to extend the previous results to an inconsistent source sampling scheme, consider a MC simulation where the desired birth particle distribution and weight is given by the functions $p'(\mathbf{x})$ and $w(\mathbf{x})$, respectively, but inconsistent importance sampling is used to actually sample the initial particle states from an alternate

distribution, $p^*(\mathbf{x})$, with associated particle weights $w^*(\mathbf{x})$. The sampled particle states can be converted to the desired weight, $w(\mathbf{x})$, by applying an unbiased splitting/rouletting process with weight adjustment factor

$$M(\mathbf{x}) = \frac{p'(\mathbf{x})}{p^*(\mathbf{x})} = \frac{w^*(\mathbf{x})}{w(\mathbf{x})}. \quad (59)$$

Recall that the splitting/rouletting process will terminate or increase the weight of particles where $M(\mathbf{x}) < 1$ and will split and decrease the weight of particles when $M(\mathbf{x}) > 1$.

Note that the analyses in the previous section assumed either all splitting or rouletting (with a constant weight-adjustment factor M) for source particles, whereas a realistic inconsistent source sampling scheme using Eq. (59) will involve both splitting and rouletting, with a phase-space-dependent weight adjustment factor. However, it is still possible to make general observations of the expected behavior(s) based on Eqs. (38) and (44) that result from the simplified analysis.

In an inconsistent sampling scheme, the splitting/rouletting process effectively redistributes the time spent on transport for each initial particle by not transporting (and recording zero response for) some particle states that are oversampled (relative to the target distribution $p'(\mathbf{x})$) and simulating multiple realizations for particle states that are undersampled. It is important to note that the expected number of transported particles (including replicates) is preserved by the splitting/rouletting process.

Previous work has noted that inconsistent source sampling is computationally inefficient because of the cost of rouletting and splitting source particles at birth [2,3]. However, the cost of the rouletting and splitting operations is typically small when compared to the cost of the transport for the redistributed particles. Thus, the computational inefficiency from inconsistent source sampling appears to be caused by the increase in the total variance of the response generated by the redistributed particles rather than by the process of weight adjustment itself.

Although the total amount of transport work is conserved between splitting and rouletting events, Eqs. (38) and (44) suggest that splitting and rouletting may have very different effects on the total variance of the response. The increase in variance due to splitting depends on the variance of the response with respect to only the sampled source distribution, whereas the increase in variance due to rouletting depends on the expected response for the system.

Thus, splitting is favorable in regions where the sampled source particle weight is close to the optimal weight [Eq. (23)], which causes the source variance term to go to zero, or for problems/regions where the source variance term is naturally small. Rouletting is favorable when the expected response for the source particle is close to zero. These observations suggest that it may be possible to improve the variance for an inconsistent source distribution by uniformly rescaling the number of source samples/initial particle weights to emphasize either splitting or rouletting as appropriate.

It is also important to note that the conditions where splitting/rouletting do not increase variance (zero source variance for splitting, and zero response for rouletting) are well aligned with the ideal source distribution for weight window variance reduction techniques such as CADIS. For example, applying splitting to particles sampled from a nearly ideal consistent source will not produce a significant increase in response variance because the source variance term is close to zero. Similarly, the variance increase for rouletting is minimized in regions with the lowest expected response, which is exactly where roulette is most likely to occur when sampling from an inconsistent source distribution.

Based on these observations and preliminary analytical and numerical results, it appears that the penalty (increased variance per unit work) associated with using an inconsistent source sampling scheme for MC simulations with weight window variance reduction is strongly dependent on the nature of the problem (e.g., the source variance term) as well as the deviation between the sampled source distribution and the optimal source distribution.

Prior to proceeding, it is worthwhile to discuss the context of the theoretical analysis presented in this paper, along with the known limitations of the analysis and opportunities for future investigations. By assuming that the weight adjustment parameter, M , is constant, the analysis presented in this work effectively decouples the variance penalty associated with splitting and rouletting from the expected weight adjustment associated with a particular inconsistent source definition. That said, the systematic and qualitative analysis of the weight adjustment variance penalty (parameterized by the factor M) establishes a solid foundation for subsequent analyses of particular inconsistent source sampling strategies.

In addition, the characterization of the variance for a general importance sampling source method in terms of the biased source distribution $p'(\mathbf{x})$ and the associated weighting function $w(\mathbf{x})$ (Section 2.3) enables a quantitative assessment of the variance penalty associated with using importance sampling with a nonoptimal biased source distribution. This is especially significant because it opens the possibility of analytically determining the effects of weight-window discretization (in space and energy) on response variance.

Ideally, the analyses presented in this paper should be extended to account for phase-space-dependent weight adjustment parameters, $M(\mathbf{x})$, which depend on the importance sampling weighting function, $w(\mathbf{x})$, as shown in Eq. (59).

Alternatively, it is worth noting that the variance of a phase-space-dependent weight-adjustment scheme can be approximated by discretizing the phase space into a set of volumes k , each with a constant weight adjustment parameter M_k . If the set of discretized phase volumes partition the entire phase space, then the combined variance of a source sample in the problem can be derived from the Law of Total Variance, giving

$$\text{Var}[r] = \sum_{k=1}^K \text{Var}[r|k]p(k) + \sum_{k=1}^K E[r|k]^2(1 - p(k))p(k) - 2 \sum_{k=1}^K \sum_{k'=1}^K E[r|k]E[r|k']p(k)p(k'), \quad (60)$$

where the variance and expectation terms on the right-hand side of Eq. (60) are conditioned on the initial source site being produced in volume k . Although cumbersome, it is possible to expand Eq. (60) using the appropriate variance expression from Section 2 for each volume k in order to yield an estimate for the overall variance of the response. Note that as the number of partitions K is increased, the variation of the weight adjustment parameter in each partition will decrease, thus improving the suitability of the variance estimates produced in Section 2.

4. Numerical results

In order to confirm the analytical results regarding the variance for importance sampling, both with and without splitting and rouletting, a series of numerical tests were conducted using a simplified system.

The test scenario uses a discrete two-region geometry where source particles are born in either region A or B, each with source probability 0.5. For particles born in region A, the response r is

normally distributed with $\mu_A = 5$ and $\sigma_A^2 = 1$. For region B, the response is normally distributed with $\mu_B = 1$ and $\sigma_B^2 = 1$. For this situation, it can be shown that the average response is $\bar{r}=3$ and the variance of the response is $\sigma_r^2=4$. A simple program was written to sample source sites and the corresponding responses using the importance sampling schemes described in Section 2 and then compute the variance based on the observed population of responses.

Fig. 1 illustrates the dependence of the transport variance, source variance, and total variance terms from Eq. (21) for importance sampling (without splitting or rouletting) as a function of the importance sampling source weight for region A, denoted w_A . Also shown is the sample variance based on 10,000 independently sampled source locations for 100 sample weight w_A values. Note that although the source variance term is zero for $w_A = 0.6$, as predicted by Eq. (23), the minimum response variance actually occurs at a slightly higher value of $w_A = 0.6287$ as a result of the interplay between the source and transport variance terms, as described in Section 2.

Figs. 2A and 2B show the response variance as a function of sampled weight w_A and splitting factor/roulette survival probability M . Both cases (splitting and rouletting) show simulations for 100 values of w_A and five different values for M , from 1.0 to 10.0 for splitting and 0.1 to 1.0 for rouletting. In each simulation, the total amount of transport work was maintained constant by adjusting the number of source samples N' such that $N' \cdot M = 10,000$. Inspection of the results shows excellent agreement with Eqs. (38) and (44), as expected.

5. Conclusions

This summary provides a novel analysis of the variance associated with importance sampling both with and without uniform splitting and rouletting. In this analysis, the law of total variance is used to demonstrate the interplay between two separate sources of variance when importance sampling schemes are used: variance due to the transport process itself and variance due to the effect of the source distribution. Results of the analysis demonstrate that both splitting and rouletting tend to increase total response variance in proportion to either the source variance term (splitting) or the expected response for the system (rouletting). These observations, in turn, can be used to predict the potential effects of an inconsistent source sampling scheme on total response variance. Based on the analysis of the variance performed for the different weight adjustment schemes, it appears that the penalty associated with inconsistent source sampling is strongly dependent on the

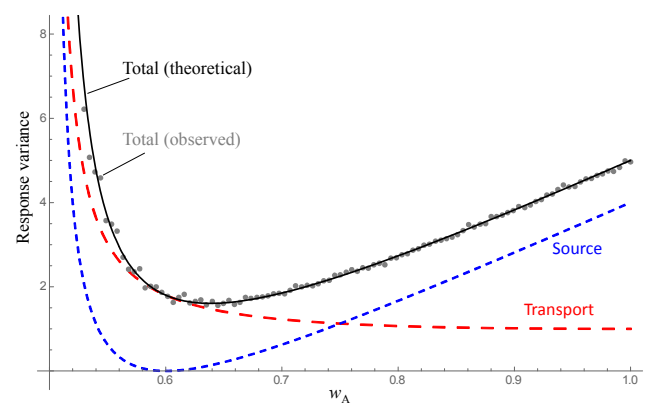


Fig. 1. Response variance as a function of importance sampling weight for region A. Theoretical results for transport, source, and total variance are shown, along with experimental results for total variance.




Fig. 2. Response variance as a function of importance sampling weight for region A for various particle split/survival factors. (A) Effect of source splitting on response variance. (B) Effect of rouletting on response variance.

inherent response variability of the problem itself as well as the difference between the source distribution used for sampling and the optimal (i.e., adjoint-driven) importance sampling source distribution. Early testing also suggests that certain classes of problems may be relatively insensitive to inconsistent source sampling schemes with moderate levels of splitting and rouletting.

Conflicts of interest

The authors affirm that they do not have any potential conflicts of interest related to this research.

Acknowledgments

The authors thank Dr. Brian Nease of the Naval Nuclear Laboratory for helpful discussions on Monte Carlo variance reduction.

References

- [1] G.I. Bell, S. Glasstone, *Nuclear Reactor Theory*, Van Nostrand and Reinhold, New York, 1967.
- [2] J.C. Wagner, A. Haghghat, Automated variance reduction of Monte Carlo shielding calculations using the discrete ordinates adjoint function, *Nucl. Sci. Eng.* 128 (1998) 186–208.
- [3] J.C. Wagner, D.E. Peplow, S.W. Mosher, FW-CADIS method for global and regional variance reduction of Monte Carlo radiation transport calculations, *Nucl. Sci. Eng.* 176 (2014) 37–57.
- [4] S.C. Wilson, D.P. Griesheimer, Rejection sampling from a spherical harmonics angular flux functional expansion for a discrete ordinates to Monte Carlo splice, *Proc. International Conference on Mathematics, Computational Methods & Reactor Physics (M&C 2009)*, CD-ROM, Saratoga Springs, NY, 2009.
- [5] X-5 Monte Carlo Team, MCNP – A General Monte Carlo N-Particle Transport Code, Version 5, LA-CP-03-0284, Los Alamos National Laboratory, 2003.
- [6] T.E. Booth, J.S. Hendricks, Importance estimation in forward Monte Carlo calculations, *Nucl. Technol. Fusion* 5 (1984) 90–100.
- [7] J. Spanier, E.M. Gelbard, *Monte Carlo Principles and Neutron Transport Problems*, Dover Publications, New York, 2008.