Eurographics Symposium on Rendering (2021)
A. Bousseau and M. McGuire (Guest Editors)

Zero-variance transmittance estimation

Eugene d'Eon   Jan Novák

NVIDIA

Abstract

We apply zero-variance theory to the Volterra integral formulation of volumetric transmittance. We solve for the guided sampling decisions in this framework that produce zero-variance ratio tracking and next-flight ratio tracking estimators. In both cases, a zero-variance estimate arises by colliding only with the null particles along the interval. For ratio tracking, this is equivalent to residual ratio tracking with a perfect control. The next-flight zero-variance estimator is of the collision type and can only produce zero-variance estimates if the random walk never terminates. In drawing these new connections, we enrich the theory of Monte Carlo transmittance estimation and provide a new rigorous path-stretching interpretation of residual ratio tracking.

CCS Concepts

• ***Computing methodologies*** → ***Reflectance modeling;***

1. Introduction

Evaluating visibility between points in a scene is a fundamental task performed by light transport simulators. If the scene contains opaque surfaces only, visibility is a binary function. If a participating medium is present, however, visibility becomes a fractional quantity that corresponds to the probability of traversing between the points without interacting with the medium. If the medium is random and the size of its particles negligible, the radiative transfer theory [Cha60] can be used to statistically model interactions with absorbing and scattering particles. The number of particles along a given ray is then a random variable, and the probability of flying along the ray from a to b without hitting particle is commonly referred to as *transmittance* [Pre65]:

$$T(a,b) = \exp\left(-\int_a^b \mu(x)dx\right), \tag{1}$$

where the extinction coefficient μ quantifies the probability density of interacting with a particle per unit-length flight. The coefficient is deterministic, non-negative and generally varies as a function of position.

Accurately estimating transmittance is essential for efficient rendering of scenes with haze, smoke, or clouds. The integral of $\mu(x)$ in Equation (1) is known in closed form only for a few, rather simple extinction functions (e.g. homogeneous or exponentially decreasing). The general-purpose approach, therefore, is to point-sample $\mu(x)$ at a number of positions x along the ray. It is non-trivial to form unbiased estimators from these point samples due to the nonlinearity of the exponential.

There are a number of different unbiased Monte Carlo estimators for transmittance, and each can be derived in a number of ways [GMH*19, JKU*20, KdPN21]. Of particular interest is the recent Volterra-integral-equation formulation [GMH*19] that draws a close connection to the random-walk solutions of Fredholm integral equations. This formulation is interesting because it views the estimation problem as a random walk over the one-dimensional ray/interval, with strong analogies to path-tracing and related Monte Carlo solutions of general transport problems. These analogies have inspired a number of new estimators and interpretations of previous estimators [GMH*19].

It is known that most random-walk solutions of Fredholm integral equations can be guided using modified sampling strategies such that every sampled random walk gives the exact answer with zero variance. The purpose of this paper is to combine this zero-variance theory with the Volterra formulation for transmittance to answer the question: *what are the zero-variance estimators for volumetric transmittance*? We derive two such estimators in the following sections and explore their properties and relationships to prior work. By expanding upon the theory of Monte Carlo transmittance estimation in this way we open the door to possible new strategies for computing low-variance visibility when rendering scenes with rich atmospheric effects, and we discuss some of these ideas at the end of the paper.

Scope. We do not present any new practical algorithms or empirical results in this article. The sole purpose of this work is to provide a novel derivation of variance-optimal estimators and to show that a certain existing estimator can be viewed as an approximation of one of these optimal estimators.

2. Background

In this section, we review the Volterra integral equation of transmittance and the zero-variance theory, which we combine in later sections to derive zero-variance last-event and collision estimators.

2.1. Volterra formulation of volumetric transmittance

The *Volterra* formulation of transmittance follows from the observation that Equation (1) satisfies the Volterra integral equation [GMH*19, Eq.(9)]

$$T(a,b) = 1 - \int_a^b \mu(x) T(x,b) \mathrm{d}x, \tag{2}$$

which expresses transmittance using a recursive rule; T now appears also on the right-hand side. A Monte Carlo estimator derived from Equation (2) randomly decides whether or not to evaluate the source term (the "1") with some probability, and (independently) whether or not to recurse to a subinterval by sampling a point $X \in (a,b)$ and then estimating $T(X,b)$ with the same procedure. Corresponding weight factors are then introduced to make the estimates unbiased. Thus, the estimator effectively performs a random walk along the interval (a,b), querying the extinction $\mu(x)$ at N locations along the interval (where N is typically Poisson-distributed) and computing an unbiased estimate from those N queries.

Georgiev et al. [GMH*19] showed that previous Monte Carlo transmittance estimators such as ratio tracking [Cra78, NSJ14] can be reinterpreted under this new integral formulation. Most estimators, like ratio tracking, are based on the concept of null collisions, where ficticious particles are introduced along the interval in such a way that they make collision sampling simpler while not changing the light transport in any way [NGHJ18]. This is achieved by specifying a *majorant* $\bar{\mu}(x)$ satisfying $\mu(x) \leq \bar{\mu}(x)$ for $a \leq x \leq b$. The extinction coefficient for the newly-introduced null particles is then $\mu_\mathrm{n}(x) = \bar{\mu}(x) - \mu(x)$. The majorant can be used as a control variate to transform Equation (2) into [GMH*19, Eq.(13)]:

$$T(a,b) = T_{\bar{\mu}}(a,b) + \int_a^b T_{\bar{\mu}}(a,x) \mu_\mathrm{n}(x) T(x,b) \mathrm{d}x. \tag{3}$$

Here we use the notation

$$T_{\bar{\mu}}(a,b) = \exp\left(- \int_a^b \bar{\mu}(z) \mathrm{d}z \right), \tag{4}$$

which is known analytically for any (a,b) by construction. The general form of a Volterra transmittance estimator is then

$$\langle T(a,b) \rangle = \underbrace{\frac{T_{\bar{\mu}}(a,b)}{P_\mathrm{src}}}_{\text{with prob. } P_\mathrm{src}} + \underbrace{\frac{\mu_\mathrm{n}(X) T_{\bar{\mu}}(a,X)}{P_\mathrm{rec}\, p_{ab}(X)} \langle T(X,b) \rangle}_{\text{with prob. } P_\mathrm{rec}}, \tag{5}$$

where

- P_src is the probability of evaluating the source term,
- P_rec is the probability of estimating the integral and recursing,
- $P_\mathrm{rec}\, p_{ab}(X)$ is the probability density for sampling a single random sample $X \in (a,b)$ drawn from some distribution with pdf $p_{ab}(X)$ defined on the interval (a,b).

A number of estimators follow from this formulation, which differ in the choice of free parameters P_src, P_rec, and p_{ab}. After recalling zero-variance theory in the next section, we review two of these estimators and derive their corresponding zero-variance forms. The key difference of our work to the prior analysis of the Volterra formulation [GMH*19] is in derivation of p_{ab} that achieves low or zero-variance estimates.

2.2. Zero-variance theory

Our core contribution is inspired by the zero-variance estimation of solutions of Fredholm integral equations [Alb53]. We briefly review this theory for the Fredholm case in order to recall the similarities with Equation (2) and also to clarify some key differences.

A Fredholm integral equation of the second kind is an equation

$$f(x) = q(x) + \int_a^b K(x,y) f(y) \mathrm{d}y \tag{6}$$

for unknown $f(x)$ given source term $q(x)$ and kernel $K(x,y)$. From the beginning of the development of the Monte Carlo method it was recognized that random walks can be used to estimate solutions to Equation (6) [SGS51, SG69]. The walks are terminated with some procedure, such as Russian roulette, and there are various different estimators that can be formed from the same walk, such as the *last-event estimator* that only contributes to the estimate on the final step of the walk, or the *collision estimator* that contributes to the estimate at each vertex.

It was later shown that there exists a theoretical *importance-sampling* of the random walk for Equation (6) such that the last-event [Kah56] or collision [CCY67] estimators give the exact answer with zero-variance [Hoo08]. In practice this requires knowledge of an importance function for the problem that follows from a solution to an adjoint transport problem (where particles flow from the detector to the source), which is harder to solve than the original problem. However, by merely approximating the importance function and using that to guide the construction of the walk, variance can be greatly reduced in practice. The rendering equation [Kaj86] and other equations of linear transport theory are of the Fredholm type and it is this integral form that Monte Carlo methods, such as path tracing, derive from [PJH16]. Adjoint importance and guiding have been used as variance reduction tools in these contexts [Chr03, VK16, HZE*19, dK20].

Our present investigation is inspired by the similarity between Equation (2) and Equation (6). We are not the first to exploit this similarity: the analogies between the two have been used to inspire derivation of new transmittance estimators [GMH*19]. However, we are the first (to the best of our knowledge) to apply zero-variance theory to Monte Carlo estimation of a Volterra integral equation. One key difference is that the Volterra random walks are unidirectional, in that the recursion reduces the problem to strictly smaller domains at each step. Invariant imbedding can yield integral formulations for transport problems with a similar property [PH00], but in general Monte Carlo estimators and zero variance theory are considered in the context of Equation (6). We found that this difference makes the derivation of the Volterra zero-variance guiding simpler and more direct.

3. The last-event estimator

We first focus on *last-event* estimators that alter the Monte Carlo score only at the last simulated event (i.e. during the last step of the recursion).

3.1. Standard ratio tracking

Consider the transmittance estimator that follows from Equation (3) where the sampled locations X follow from a particle flying from point a to point b and colliding with the sum of real and null particles $\mu(X)+\mu_{\text{n}}(X)=\bar{\mu}(X)$, i.e. $p(X)=\bar{\mu}(X)T_{\bar{\mu}}(a,X)$ (these estimators follow the track of a particle as it traverses the interval and are also known as *tracking* estimators). Consider further that the estimator evaluates the source term in Equation (3) only when the particle exceeds point b, i.e. $P_{\text{src}}=P\{X>b\}=T_{\bar{\mu}}(a,b)$, and otherwise the second term is estimated at the sampled location X; the probability density of sampling X is $P_{\text{rec}}\,p_{ab}(X)=p(X)=\bar{\mu}(X)T_{\bar{\mu}}(a,X)$. Georgiev et al. [GMH*19] noted that this leads to an estimator known as *weighted tracking on a line* [Cra78], in graphics often referred to as *ratio tracking* [NSJ14]. The estimator steps along the line computing the product of local weights (ratios of null to total particle densities) $\frac{\mu_{\text{n}}(X)}{\bar{\mu}(X)}$. When $X>b$, we evaluate the first term, which simplifies to 1 with the aforementioned PDF $p(X)$, terminate the walk, and score the product of ratios.

The variance of this estimator will generally depend on the variance of the ratios and the variance of the number of random steps that the estimator performs.

3.2. Zero-variance ratio tracking

We will now show how to eliminate variance in the ratio tracking estimator. Standard ratio tracking samples the distance X using the majorant extinction: $p(X)=\bar{\mu}(X)T_{\bar{\mu}}(a,X)$. Since we are free to choose the probability density of sampling distances, we might as well set it such that a zero-variance estimate is achieved. To eliminate all variance, we need the estimator to return the true transmittance for the interval (a,b) at any point during the recursive invocation—a constraint that we impose over the free parameters.

In ratio tracking, the free parameters P_{src}, P_{rec}, and p_{ab} are linked and derived from the sampling density p. We therefore need to impose the aforementioned constraint on one of the free parameters only and the others will follow.

Focusing on the first (source) term in Equation (5), the constraint

$$\frac{T_{\bar{\mu}}(a,b)}{P_{\text{src}}}=T(a,b) \tag{7}$$

implies

$$P_{\text{src}}=\exp\left(-\int_a^b \bar{\mu}(z)-\mu(z)\mathrm{d}z\right) \tag{8}$$

$$=\exp\left(-\int_a^b \mu_{\text{n}}(z)\mathrm{d}z\right)=T_{\mu_{\text{n}}}(a,b)\,. \tag{9}$$

Since $1-P_{\text{src}}$ is the CDF of $p(X)$, the constraint is satisfied with distance sampling density

$$\boxed{p_{\text{zv}}(X)=\mu_{\text{n}}(X)\exp\left(-\int_a^X \mu_{\text{n}}(z)\mathrm{d}z\right).} \tag{10}$$

The density p_{zv} is the free-flight distribution for a particle that collides with *only the null particles* of the medium. As such, ratio tracking should sample collisions with the null particles only to yield zero variance.

The density p_{zv} can also be derived by imposing the constraint on the second term of Equation (5):

$$\frac{\mu_{\text{n}}(X)T_{\bar{\mu}}(a,X)}{P_{\text{rec}}\,p_{ab}(X)}=T(a,X)\,, \tag{11}$$

and noting that in ratio tracking $P_{\text{rec}}\,p_{ab}(X)=p(X)$.

The same result can be derived yet a third way, by using the equivalence of the ratio-tracking estimator to a single-term power series estimator, as show by Georgiev et al. The general expression for such an estimator can be written [GMH*19, Eq.(26)]

$$\langle T(a,b)\rangle=\frac{e^{-\bar{\tau}}}{k!P(k)}\prod_{i=1}^{k}\frac{\mu_{\text{n}}(X_i)}{p(X_i)} \tag{12}$$

where $P(k)$ is the probability that k collisions are sampled, each collision $X_i\in(a,b)$ independently sampled with PDF $p(X_i)$. To apply our zero-variance constraint to Equation (12), the product $\prod_{i=1}^{k}\frac{\mu_{\text{n}}(X_i)}{p(X_i)}$ must be zero-variance, regardless of k. This clearly can only happen when $p(X_i)$ is proportional to $\mu_{\text{n}}(X_i)$, which is when each X_i is sampled in proportion to $\mu_{\text{n}}(X_i)$. This is exactly what delta tracking does when only colliding with the null particles (by the independence property of Poisson processes), and produces a Poisson-distributed $P(k)\sim\text{Po}(\tau_{\text{n}})$ number of points X_i. The product reduces to $\prod_{i=1}^{k}\frac{\mu_{\text{n}}(X_i)}{p(X_i)}=\tau_{\text{n}}^k$ and Equation (12) reduces to $e^{-\tau}$ for all k. Therefore, we have also derived the unique zero-variance, single-term power series estimator.

Note that the zero-variance estimator for a Fredholm integral equation can be derived in the same fashion: by requiring that the first sampling decision results in a zero-variance result [dK20]. However, this invariably leads to a product distribution with an unknown normalization factor (which happens to be the quantity we set out to estimate). One term in this product distribution is the unknown importance function, and so there are two intractable unknowns in the final estimator. While neither of these appear here in the Volterra case, we have nevertheless arrived at an intractable result: we do not know Equation (10).

3.2.1. Discussion

The zero-variance density p_{zv} does not immediately lead to a practical numerical algorithm; this is a common drawback of zero-variance schemes. To illustrate the reasons, we rewrite Equation (5) replacing the generic parameters with expressions that arise when p_{zv} is used:

$$\langle T(a,b)\rangle_{\text{zv}}=\begin{cases}\dfrac{T_{\bar{\mu}}(a,b)}{T_{\mu_{\text{n}}}(a,b)} & \text{if } X>b\text{, and}\\[2ex] \dfrac{\mu_{\text{n}}(X)T_{\bar{\mu}}(a,X)}{p_{\text{zv}}(X)}\langle T(X,b)\rangle & \text{otherwise.}\end{cases} \tag{13}$$

In the first case, the ratio of transmittances simplifies to the true value of transmittance: $T_{\bar{\mu}}/T_{\mu_n} = T$. This is the constraint that we imposed. It makes the estimator zero-variance in theory, but it also makes the evaluation intractable for general heterogeneous volumes in practice.

The second case features the same problem. Suppose we use a majorant $\bar{\mu}(x)$ that permits evaluating $T_{\bar{\mu}}(a,x)$ in closed form, e.g. a constant function. Since the majorant is relatively simple, the null extinction $\mu_n(x)$ will feature variations that are comparable to those in the real extinction $\mu(x)$, hence we cannot sample the distance analytically. One might exercise the idea of utilizing delta tracking to sample distances in the null medium. This would yield the desired sample density p_{zv}, but as this density cannot be evaluated when delta tracking is used, our only hope is that it cancels out with the terms in the numerator (up to an easy-to-compute factor). Since the fraction simplifies to $T(a,X)$, we again end up with a quantity that is intractable to evaluate.

This is a universal property of zero-variance schemes—there is no point in using them when they are truly zero-variance because they require solving a harder problem than we set out to solve in the first place. However, they become valuable in showing exactly how a given estimator should be perturbed so that its variance is lowered, and what necessary constraints need to be met to enable tractable evaluation. In this specific instance, we shall strive to sample only null collisions when generating distances, but not at the cost of making the evaluation of other terms intractable.

There also remains the question of cost, which can be arbitrarily high. Note that for any strict majorant, $p_{zv}(x)$ is well defined. As the majorant is made closer to the control, the cost goes down, fewer points are sampled on average along the interval (there is less null mass between a and b), but the estimator remains zero-variance. So unlike the analogous Fredholm random walk, there is *not* a unique last-event zero-variance estimator. We should strive for variants that reduce not only variance but also the cost in practice.

3.3. Approximate zero-variance ratio tracking

In order to create a tractable low-variance estimator, we define the *approximate* null extinction along the (a,b) segment to be $\tilde{\mu}_n(x) = \bar{\mu}(x) - \tilde{\mu}(x)$; here, $\tilde{\mu}(x)$ can be interpreted as an approximation of the real extinction function $\mu(x)$.

The approximate null extinction $\tilde{\mu}_n(x)$ can be any non-negative function for $x > a$ provided that $\tilde{\mu}_n(x) > 0$ whenever $\mu_n(x) > 0$. We also demand that the transmittance $T_{\tilde{\mu}_n}(a,x) = \exp\left(-\int_a^x \tilde{\mu}_n(z)\mathrm{d}z\right)$ can be evaluated exactly, and $\int_b^\infty \tilde{\mu}_n(x)dx = \infty$. These conditions merely ensure that we can generate finite distance samples and evaluate the transmittance in this new approximate null medium.

The approximately zero-variance, ratio-tracking estimator that samples distances using $p_{azv}(X) = \tilde{\mu}_n(X)T_{\tilde{\mu}_n}(a,X)$ reads:

$$\langle T(a,b)\rangle_{azv} = \begin{cases} \dfrac{T_{\bar{\mu}}(a,b)}{T_{\tilde{\mu}_n}(a,b)} & \text{if } X > b, \text{ and} \\ \dfrac{\mu_n(X)T_{\bar{\mu}}(a,X)}{p_{azv}(X)}\langle T(X,b)\rangle & \text{otherwise.} \end{cases} \tag{14}$$

Comparing this estimator to Equation (13), we merely replaced the sampling PDF and the null-medium transmittance with their corresponding counterparts derived from the approximate null extinction $\tilde{\mu}_n(x)$. The benefit of this substitution is that all terms on the right-hand side can now be easily evaluated, albeit at the cost of losing the zero-variance property.

3.3.1. Correspondences with residual ratio tracking

The approximate zero-variance estimator in Equation (14) corresponds *identically* to the *residual ratio tracking* (RRT) estimator [NSJ14] (which is equivalent to the Poisson estimator [BPRF06, KdPN21]), the only difference being the nomenclature. We will now outline these correspondences.

In this article, we utilized the approximate real extinction $\tilde{\mu}(x)$, which maps to the "control extinction"—denoted $\mu_c(x)$—in RRT; in both cases this quantity serves as a close approximation to the real extinction. We also introduced the approximate null extinction $\tilde{\mu}_n(x)$, which corresponds to the "majorant of the residual medium"—denoted $\bar{\mu}_r(x)$—in RRT; both articles propose to generate distances according to this quantity.

As such, we have not presented a new algorithm for estimating transmittance, but merely a new derivation of an existing estimator. Nevertheless, there are two important insights gained from this derivation:

- It is evident that one can minimize variance by sampling collisions in the null medium only; this has not been obvious from prior works.
- The introduction of the residual medium [NSJ14] serves as a step towards introducing the *majorant* of the residual medium. While being perhaps didactic, we show that the concept of the residual medium is unnecessary as we only need its majorant (the approximate null extinction $\tilde{\mu}_n(x)$ in our terminology) to draw the distance samples.

Our work complements the original RRT article by providing a more direct, and arguably more intuitive, way of deriving a variance minimizing transmittance estimator. Additionally, we note close analogies to path-stretching methods first introduced in approximating zero-variance walks in neutron transport—specifically the *exponential transform* that results from sampling free-flight distances in homogeneous media using an altered constant extinction coefficient [NGHJ18, dK20]: this is exactly what residual ratio tracking does when the majorant and control are both constant over the interval (a,b).

Zero variance theory is typically applied by deriving the theoretical result and then testing an approximate importance function to measure the increase in efficiency as the error in the approximation is reduced. Due to the correspondence with RRT, these studies have already been performed and we will not repeat them here. When the majorant and the control are restricted to be constants over the interval, the optimum value for the control is known [BPRF06, KdPN21], which follows from the simple expression for the exact variance of the RRT estimator. The efficiency of non-constant controls has also been measured as a function of the degree of piecewise-polynomial approximations and their resolution [NSJ14, Fig.(13)].

4. The collision estimator

In this section, we discuss the application of zero-variance theory to next-flight estimators. Specifically, we consider the *weighted next-flight (WNF) estimator* by Cramer [Cra78] (in graphics sometimes dubbed the *next-flight ratio tracking estimator*). Georgiev et al. [GMH*19] noted that this estimator emerges from Equation (5) when P_{src} is set to 1 (i.e. the source term is always evaluated) and the probability density of recursing at point x is kept the same as in standard ratio tracking, again using free-flight distances arising from collisions with the total medium (i.e. $P_{\text{rec}}\, p_{ab}(X) = \bar{\mu}(X) T_{\bar{\mu}}(a,X)$).

For each sampled distance point, the WNF estimator computes a score consisting of the transmittance along the remaining segment considering the combined medium $T_{\bar{\mu}}(X,b)$, and scales it by the product of $\frac{\mu_{\text{n}}(X)}{\bar{\mu}(X)}$ weights from all preceding sampled locations. These scores are added together to form the transmittance estimate. In this way, instead of scoring a single weight at the end when the source term is finally evaluated (as in last-event estimators), here a contribution to the transmittance is made at each collision. This maps exactly to the collision estimator of zero-variance theory.

In order to derive the zero-variance version of this estimator, we constrain it to always return the correct transmittance value, i.e. $\langle T(x,y)\rangle := T(x,y)$. The sampling PDF $p_{\text{zv}}(X)$ can be derived directly from Equation (5) where all estimators are replaced by the actual transmittance functions (and $P_{\text{src}} = 1$):

$$T(a,b) = T_{\bar{\mu}}(a,b) + \frac{T_{\bar{\mu}}(a,X)\mu_{\text{n}}(X)}{p_{\text{zv}}(X)} T(X,b). \tag{15}$$

The PDF

$$p_{\text{zv}}(X) = \frac{T_{\bar{\mu}}(a,X)\mu_{\text{n}}(X)T(X,b)}{T(a,b) - T_{\bar{\mu}}(a,b)} \tag{16}$$

$$= \frac{T(a,b)\, T_{\mu_{\text{n}}}(a,X)\, \mu_{\text{n}}(X)}{T(a,b)\left(1 - \frac{T_{\bar{\mu}}(a,b)}{T(a,b)}\right)} \tag{17}$$

and can be further simplified to

$$\boxed{p_{\text{zv}}(X) = \frac{\mu_{\text{n}}(X)\, T_{\mu_{\text{n}}}(a,X)}{1 - T_{\mu_{\text{n}}}(a,b)}.} \tag{18}$$

The zero-variance PDF is valid (i.e. it integrates to unity) *but for the range* $[a,b]$. The numerator corresponds to the probability density of null collisions along a half line $[a,\infty]$. The denominator corresponds to the integral of this density over the range $[a,b]$. The fraction thus equals the probability density of null collisions within $[a,b]$, and so again we find that zero-variance is attained when the distance-sampling guiding is governed by a collision process that only considers the null particles along the interval. The truly zero-variance result is intractable, but an approximate density can be used to approach the limit of zero variance as we saw for ratio tracking.

There is, however, an important difference here to the ratio tracking result. In order to yield a zero-variance estimation with the next-flight ratio-tracking estimator, we *must* evaluate the second term. In other words, the recursion never stops ($P_{\text{rec}} = 1$). The zero-variance estimator thus incurs an infinite cost. This is consistent with zero-variance theory: it is known that the collision estimator requires an unterminated random walk to achieve zero variance [Hoo08]. To create a practical, approximately zero-variance estimator, one would therefore need to trade between zero variance and infinite cost, for example by returning 1 when the remaining interval to be estimated is below some threshold. Alternatively, a hybrid approach that eventually switches to a zero-variance ratio tracker would terminate the estimator and maintain zero-variance. We leave the exploration of such recipes to future work.

In the non-guided case, ratio tracking tends to outperform next-flight ratio tracking [Cra78, GMH*19], but with some exceptions. We would expect this general trend to hold in the case that both are guided towards zero variance with the same approximate null density. One interesting feature of the zero-variance next-flight estimator, compared to the original, is that the zero-variance free-flight PDFs $p_{\text{zv}}(x)$ are *correlated* (they increase by a constant as the remaining interval narrows). This correlation is missing from the original estimator and would seem to be required in order to achieve low variance results and is perhaps why next-flight tends to underperform relative to ratio tracking.

For the last-event estimator, we found that the power-series formulation immediately led to the same zero-variance derivation. It is, however, not immediately obvious from the corresponding equation for a truncated power-series estimator [GMH*19, Eq.(28)]

$$\langle T(a,b)\rangle = e^{-\bar{\tau}} \sum_{j=0}^{k} \frac{\prod_{i=1}^{j} \frac{\mu_{\text{n}}(X_i)}{p(X_i)}}{j!\left(1 - \sum_{m=0}^{j-1} P(m)\right)} \tag{19}$$

that X_i should be correlated, or even that k must be ∞ to yield a zero-variance result. Whether there are additional zero-variance power-series estimators we believe to remain an open question.

5. Conclusion and future work

We have shown that zero-variance theory can be applied to transmittance estimation and we have derived two zero-variance transmittance estimators. Both share the common feature that collisions along the interval should be sampled by considering only the null particles. By using an approximate description of the extinction coefficient along an interval to guide distance sampling, both ratio tracking and next-flight ratio tracking approach zero variance. We noted that this corresponds to residual ratio tracking and a related residual next-flight estimator that requires an infinite cost to achieve zero variance. We have also established that our zero-variance ratio tracking estimator is the unique zero-variance single-term power-series estimator for transmittance.

By reformulating variance reduction in the path-stretching framework of zero-variance theory, instead of the control/residual collision picture of prior theory, there are a number of possible online variance-reduction methods that become more directly applicable to transmittance estimation. For example, when the distance-sampling PDFs $p(x,\omega)$ at any position and direction in the volume are drawn from parametric distributions (such as exponential or Gamma), the spatial and directional parameters of these distributions could be directly informed by heuristics, online/neural predictions or possibly utilize spatio-temporal reservoir schemes [BWP*20], etc. We feel that these are interesting directions for future work.

References

[Alb53] ALBERT G.: *A general theory of stochastic estimates of the Neumann series for the solutions of certain Fredholm integral equations and related series*. Tech. Rep. ORNL-1508, Oak Ridge National Lab., Tenn., 1953. URL: https://www.osti.gov/servlets/purl/4427633. 2

[BPRF06] BESKOS A., PAPASPILIOPOULOS O., ROBERTS G. O., FEARNHEAD P.: Exact and computationally efficient likelihood-based estimation for discretely observed diffusion processes (with discussion). *Journal of the Royal Statistical Society: Series B (Statistical Methodology) 68*, 3 (2006), 333–382. URL: https://doi.org/10.1111/j.1467-9868.2006.00552.x. 4

[BWP*20] BITTERLI B., WYMAN C., PHARR M., SHIRLEY P., LEFOHN A., JAROSZ W.: Spatiotemporal reservoir resampling for real-time ray tracing with dynamic direct lighting. *ACM Transactions on Graphics (TOG) 39*, 4 (2020), 148–1. URL: https://doi.org/10.1145/3386569.3392481. 5

[CCY67] COVEYOU R., CAIN V., YOST K.: Adjoint and Importance in Monte Carlo Application. *Nucl. Sci. Eng. 27*, 1 (1967), 219–234. URL: https://doi.org/10.13182/NSE67-A18262. 2

[Cha60] CHANDRASEKHAR S.: *Radiative Transfer*. Dover, 1960. 1

[Chr03] CHRISTENSEN P. H.: Adjoints and importance in rendering: An overview. *IEEE Transactions on Visualization and Computer Graphics 9*, 3 (2003), 329–340. URL: https://doi.org/10.1109/TVCG.2003.1207441. 2

[Cra78] CRAMER S.: Application of the fictitious scattering radiation transport model for deep-penetration Monte Carlo calculations. *Nuclear Science and Engineering 65*, 2 (1978), 237–253. URL: https://doi.org/10.13182/NSE78-A27154. 2, 3, 5

[dK20] D'EON E., KRIVÁNEK J.: Zero-variance theory for efficient subsurface scattering. *ACM SIGGRAPH 2020 Courses: Advances in Monte Carlo rendering: the legacy of Jaroslav Křivánek*, 3 (2020), 1–366. URL: https://doi.org/10.1145/3388769.3407458. 2, 3, 4

[GMH*19] GEORGIEV I., MISSO Z., HACHISUKA T., NOWROUZEZAHRAI D., KŘIVÁNEK J., JAROSZ W.: Integral formulations of volumetric transmittance. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia) 38*, 6 (Nov. 2019). URL: https://doi.org/10/dffn, doi:10/dffn. 1, 2, 3, 5

[Hoo08] HOOGENBOOM E.: Zero-variance monte carlo schemes revisited. *Nucl. Sci. Eng. 160*, 1 (2008), 1–22. URL: https://doi.org/10.13182/NSE160-01. 2, 5

[HZE*19] HERHOLZ S., ZHAO Y., ELEK O., NOWROUZEZAHRAI D., LENSCH H. P., KŘIVÁNEK J.: Volume path guiding based on zero-variance random walk theory. *ACM Transactions on Graphics (TOG) 38*, 3 (2019), 1–19. URL: https://doi.org/10.1145/3230635. 2

[JKU*20] JONSSON D., KRONANDER J., UNGER J., SCHON T. B., WRENNINGE M.: Direct Transmittance Estimation in Heterogeneous Participating Media Using Approximated Taylor Expansions. *IEEE Transactions on Visualization and Computer Graphics* (2020). doi:10.1109/TVCG.2020.3035516. 1

[Kah56] KAHN H.: Applications of monte carlo. 2

[Kaj86] KAJIYA J. T.: The rendering equation. In *Proceedings of ACM SIGGRAPH* (1986), pp. 143–150. URL: https://doi.org/10.1145/15922.15902. 2

[KdPN21] KETTUNEN M., D'EON E., PANTALEONI J., NOVÁK J.: An unbiased ray-marching transmittance estimator. *ACM Transactions on Graphics (TOG)* (2021). URL: https://doi.org/10.1145/3450626.3459937. 1, 4

[NGHJ18] NOVÁK J., GEORGIEV I., HANIKA J., JAROSZ W.: Monte carlo methods for volumetric light transport simulation. In *Computer Graphics Forum* (2018), vol. 37, Wiley Online Library, pp. 551–576. URL: https://doi.org/10.1111/cgf.13383. 2, 4

[NSJ14] NOVÁK J., SELLE A., JAROSZ W.: Residual ratio tracking for estimating attenuation in participating media. *ACM Transactions on Graphics (TOG) 33*, 6 (2014), 179. URL: https://doi.org/10.1145/2661229.2661292. 2, 3, 4

[PH00] PHARR M., HANRAHAN P.: Monte carlo evaluation of non-linear scattering equations for subsurface reflection. In *Proceedings of ACM SIGGRAPH 2000* (2000), pp. 75–84. URL: https://doi.org/10.1145/344779.344824. 2

[PJH16] PHARR M., JAKOB W., HUMPHREYS G.: *Physically based rendering: From theory to implementation*. Morgan Kaufmann, 2016. 2

[Pre65] PREISENDORFER R. W.: *Radiative transfer on discrete spaces*, vol. 1. Pergamon Press Oxford, 1965. 1

[SG69] SPANIER J., GELBARD E.: *Monte Carlo principles and neutron transport problems*. Addison-Wesley Pub. Co., 1969. 2

[SGS51] SPINRAD B., GOERTZEL G., SNYDER W.: An alignment chart for Monte Carlo solution of the transport problem. *Nat. Bur. Stand. Appl. Math. Ser 12* (1951), 4–5. 2

[VK16] VORBA J., KŘIVÁNEK J.: Adjoint-driven Russian roulette and splitting in light transport simulation. *ACM Transactions on Graphics (TOG) 35*, 4 (2016), 1–11. URL: https://doi.org/10.1145/2897824.2925912. 2