

Residual Ratio Tracking for Estimating Attenuation in Participating Media (supplementary material #2)

Jan Novák^{1,2} Andrew Selle¹ Wojciech Jarosz²

¹Walt Disney Animation Studios ²Disney Research Zürich

26.10.2014

1 Non-bounding “Majorants”

In this supplementary document, we revisit the issue of non-bounding “majorants”. In Section 3.2.2 of the paper, we incorrectly stated that negative multiplicands, which may arise with non-bounding majorants, prevent convergence of the residual ratio tracking estimator in Equation (5). It turns out that the estimator remains unbiased even with negative multiplicands. In the following, we use the recently presented integral formulation of null-collision algorithms [Galtier et al. 2013] to reason about the convergence of our (residual) ratio tracking with non-bounding “majorants”.

The original delta tracking requires a *bounding* majorant. If this is violated, the “probabilities” (defined w.r.t. the majorant) of colliding with a real particle or a fictitious particle become greater than 1 or negative, respectively.¹ Fortunately, the choice of these probabilities is not constrained; they can be set to arbitrary values provided that these still represent valid probabilities and the relative concentrations of real and fictitious particles are factored into an additional *weight* carried by the random walk. As Galtier et al. [2013] point out, using the relative concentrations as the termination and the continuation probabilities is just means to intuitively explain the algorithm and reason about its correctness. The authors present a general mathematical framework that shows how to compute the weights. They also propose a new definition for the probabilities that preserves convergence even with non-bounding “majorants”.

Our ratio tracking can be seen as a variant of the original delta tracking algorithm with specifically designed termination and continuation probabilities and an additional weight. We set the termination probability to 0 for all tentative collisions that occur *before* reaching the desired distance d , and 1 otherwise. The relative concentrations of real particles are factored into a weight that the random walk scores when reaching d . Since our weights correspond to the formalization proposed by Galtier et al. [2013] we can claim unbiasedness of our algorithms even in cases when the “majorant” does not bound the (residual) extinction coefficient.

2 Numerical Verification

In this section, we provide an additional numerical analysis of variance of the residual ratio tracking estimator. The weight that the estimator scores when reaching distance d reads:

$$\langle T_r(d) \rangle_{\text{RR}} = \prod_{i=1}^K \left(1 - \frac{\mu(x_i) - \mu_c}{\bar{\mu}_r} \right), \quad (1)$$

where x_i is the i -th tentative collision point and μ_c and the $\bar{\mu}_r$ are the control extinction coefficient and the majorant residual extinction coefficient, respectively (please see the first supplementary document for a complete definition of the algorithm). Our goal is to investigate how the estimator behaves when $\bar{\mu}_r$ does not bound

$\mu_r(x)$. Since $\bar{\mu}_r$ may no longer bound $\mu_r(x)$ we refrain from using the term “majorant” and refer to $\bar{\mu}_r$ as the *free-path-sampling coefficient* (its only purpose is to generate tentative free paths).

Figures 1 and 2 show a numerical analysis of the expected value (columns (b) and (e)) and the variance (columns (c) and (f)) of the residual ratio tracking estimator for different values of μ_c and $\bar{\mu}_r$. The experiment uses a similar setup as in Figure 9 in the paper: we render an ortho-view of a unit cube filled with an absorbing medium lit by an area light source placed behind the cube. The medium changes along the z -axis only, i.e. the extinction function is identical along all camera rays. When computing the transmittance and variance images, we linearly vary the control extinction coefficient μ_c along the vertical axis of the image and the free-path sampling coefficient $\bar{\mu}_r$ along the horizontal axis. Each pixel thus corresponds to a different $\{\mu_c, \bar{\mu}_r\}$ configuration of the estimator. Each figure shows a number of different extinction-function profiles (rows). These are scaled to make the medium absorb 75% of the light (i.e. $T = 25\%$) in the three left-most columns. The extinction profiles in the three right-most columns are scaled to make the medium absorb 95% of the light (i.e. $T = 5\%$).

The yellow poly-lines mark configurations where the free-path-sampling coefficient $\bar{\mu}_r$ is the lowest but still bounding the absolute value of the residual extinction (i.e. $\bar{\mu}_r$ corresponds to the smallest possible majorant). In all configurations to the left of the yellow poly-line $\bar{\mu}_r$ does not bound μ_r .

The expected value of the transmittance (columns (b) and (e)) is the same for all configurations; we attribute the remaining difference between pixels to the variance of the estimator—shown in columns (c) and (f) as false-colored *log* variance—which is quite high for certain configurations. Our numerical analysis shares similarities with the one presented by Galtier et al. [2013]. Their weighted, delta tracking based estimator converges even with non-bounding “majorants”, but the variance in such cases increases rapidly. Our conclusions are similar. While the fact that the algorithm can handle free-path-sampling coefficients that do not bound the (residual) extinction function is certainly beneficial—it removes the burden of finding strictly bounding majorants—the significant increase in variance favors free-path-sampling coefficients that bound, or nearly bound, the (residual) extinction coefficient. Please note that the false-colored plots show the logarithm of the variance.

References

GALTIER, M., BLANCO, S., CALIOT, C., COUSTET, C., DAUCHET, J., HAFI, M. E., EYMET, V., FOURNIER, R., GAUTRAIS, J., KHUONG, A., PIAUD, B., AND TERRE, G. 2013. Integral formulation of null-collision Monte Carlo algorithms. *Journal of Quantitative Spectroscopy and Radiative Transfer* 125 (Apr.), 57–68.

¹Such quantities can no longer be considered valid probabilities, hence the quotation marks.

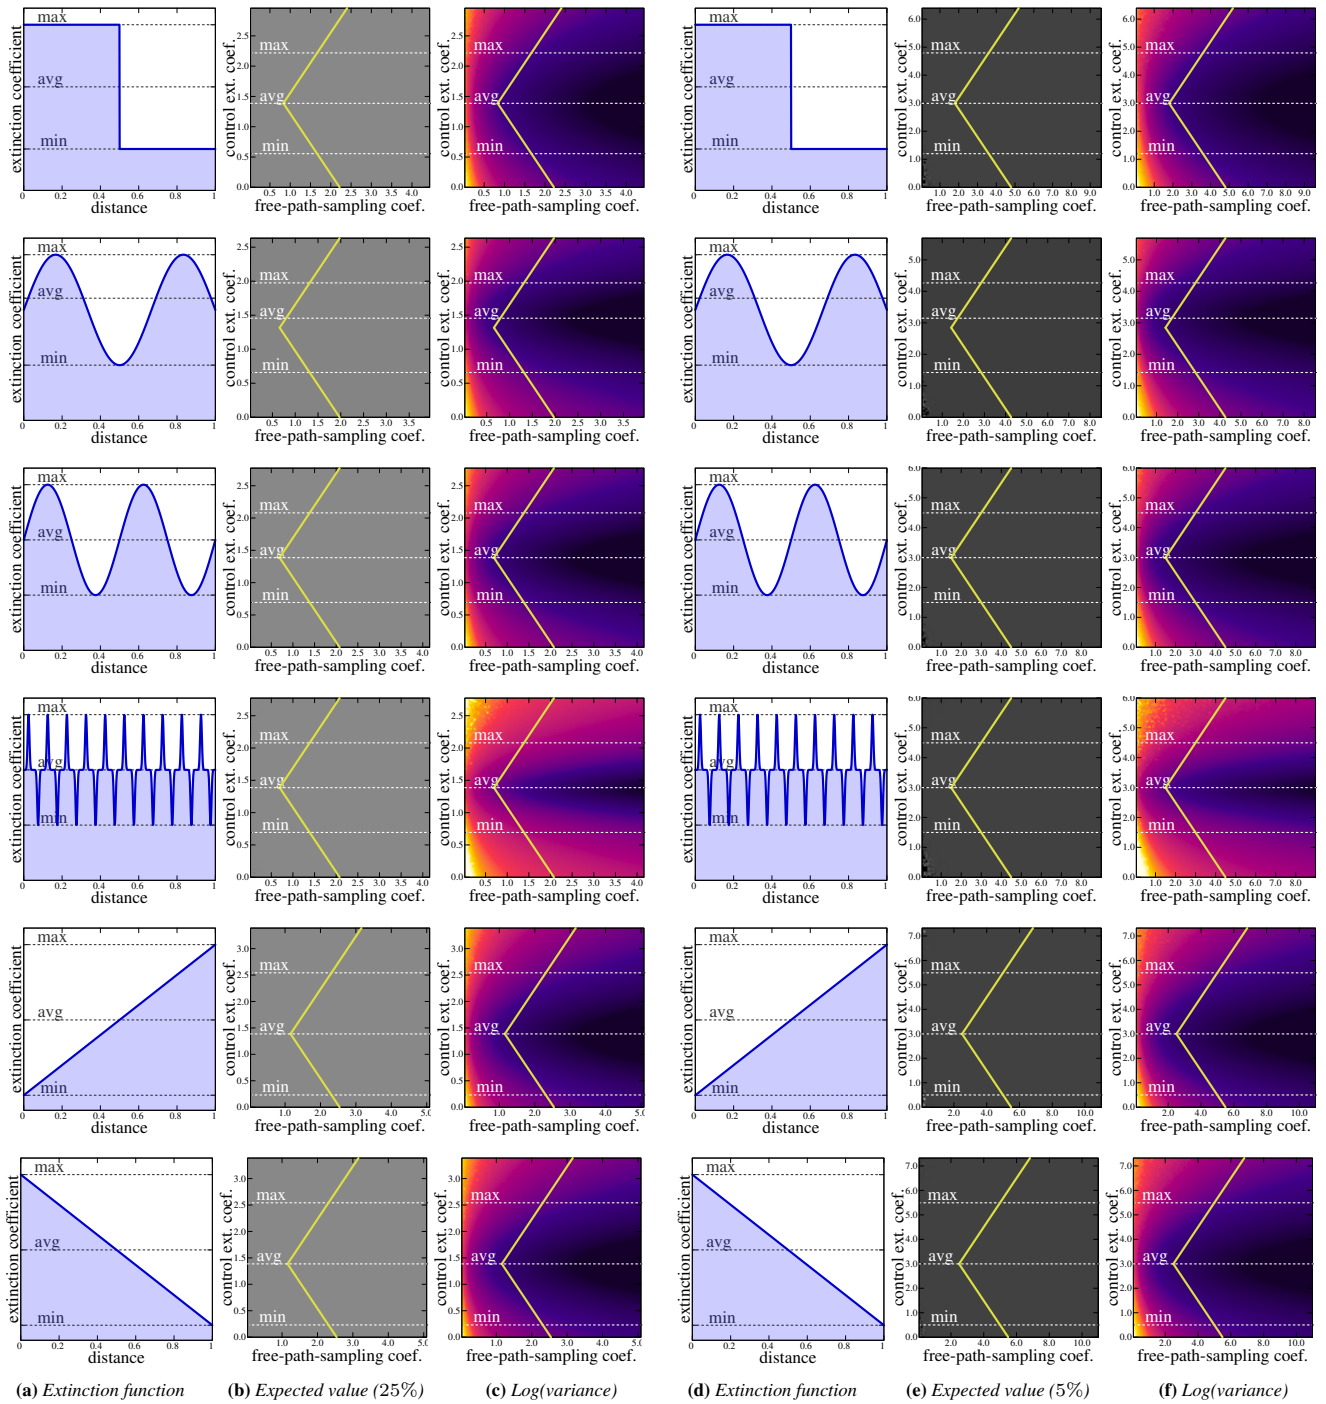


Figure 1: The expected value (columns (b) and (e)) and log variance (columns (c) and (f)) of the residual ratio tracking-based estimator for several example extinction profiles (rows) scaled to yield 25% and 5% transmittance.




Figure 2: The expected value (columns (b) and (e)) and log variance (columns (c) and (f)) of the residual ratio tracking-based estimator for several example extinction profiles (rows) scaled to yield 25% and 5% transmittance.