

Efficient Unbiased Rendering of Thin Participating Media

Ryusuke Villemin, Magnus Wrenninge, Julian Fong
Pixar Animation Studios




Figure 1. Noise bank rendered with and without our distance normalization and probability biasing methods. Inset shows magnified view.

Abstract

In recent years, path tracing has become the dominant image synthesis technique for production rendering. Unbiased methods for volume integration followed, and techniques such as delta tracking, ratio tracking, and spectral decomposition tracking are all in active use; this paper is focused on optimizing the underlying mechanics of these techniques. We present a method for reducing the number of calls to the random number generator and show how modifications to the distance sampling strategy and interaction probabilities can help reduce variance when rendering thin homogeneous and heterogeneous volumes. Our methods are implemented in version 21.7 of Pixar’s RenderMan software.

1. Introduction

Along with the film industry’s embrace of path tracing as its primary image synthesis method, volume rendering has gone from using ad-hoc lighting models and integration by quadrature (ray marching [Perlin and Hoffert 1989]) to using physically-based light transport models and unbiased integration techniques.

The move to unbiased volume integrators has been beneficial in many ways. These fit well within the path tracing framework and can efficiently evaluate even high-order multiple scattering events. Whereas ray marching depends on a (often user-specified) step length that causes bias, tracking integrators rely on control signals within the volume to guide their sampling and are able to do so in an unbiased manner. Furthermore, the control signals that drive tracking integrators can be computed automatically by either the 3D authoring application or the renderer, and they are less susceptible to user error. However, tracking integrators have certain drawbacks. They generally distribute samples according to *transmittance*, which is proportional, but not always directly related, to the radiant flux in the volume. For example, atmospheric haze due to water droplets generally has a very minor influence on transmittance, but when looking in the direction of the sun, the Mie scattering profile renders the haze orders of magnitude brighter than in other directions, even though the accumulated extinction at a given distance is identical in each case.

In this paper we describe a set of techniques that help make tracking integration more efficient in a production rendering context, with applications particularly around rendering of thin volumes and volumes with holdout objects that restrict the integration domain. We first describe a method for re-normalizing distance sampling in homogeneous volumes, providing control over the probability that a given sample will stop within a volume that is decoupled from its actual optical thickness. We then extend this notion to piecewise constant volumes by performing tracking in unit-density optical depth space, which enables a straightforward optimization of random number consumption. Finally, we introduce a probability biasing method that provides sampling control in heterogeneous volumes.

1.1. Previous Work

In production volume rendering, the transition from ray marching to physically-based path tracing has been gradual. Practitioners had long been aware of the bias inherent in quadrature-based techniques, but it was often assumed that unbiased integration methods [Woodcock et al. 1965; Skullerud 1968] that require a conservative bound $\hat{\mu}$ on extinction would be impractical in production scenes with widely varying, high-frequency extinction functions. As recently as 2011, the state-of-the-art systems, described in both SIGGRAPH courses [Wrenninge et al. 2011; Wrenninge and Bin Zafar 2011] and later in book form [Wrenninge 2012], still relied on ray marching as well as pre-computation of lighting effects using deep shadow maps [Lokovic and Veach 2000].

As an introduction to the current state of the art of volumetric light transport, we refer to the STAR report by Novák et al. [Novák et al. 2018], which provides a thorough description of the background necessary for our work, and we attempt to follow its terminology throughout.

In addition to the tracking methods themselves, work has also been done on spatial acceleration structures that provide fast lookups of extinction bounds. Yue et al. [Yue et al. 2010; Yue et al. 2011] as well as Laszlo et al. [Szirmay-Kalos et al. 2011] showed that delta tracking could be implemented as a piecewise integration of intervals, each with a locally conservative $\hat{\mu}$. Although the methods lacked treatment of features such as ray footprints and motion blur that are needed for production rendering, they highlighted the importance of designing acceleration data structures that provide well-suited intervals to a tracking-based integrator. In a SIGGRAPH course, Wrenninge [Fong et al. 2017] subsequently presented a generalized approach called *volume aggregates*, which allow for handling of ray differentials, and motion blur, as well as a large numbers of individual, overlapping volumes. We utilize the volume aggregate structure in this work in order to provide the piecewise integration intervals along with estimates of the majorant $\hat{\mu}$ and the minorant $\check{\mu}$, but other similar techniques and data structures could also be used [Szirmay-Kalos et al. 2017]. In the same course, Fong showed that sampling of homogeneous volumes could be optimized by re-normalizing the distance range (Eq. 28, p. 35.).

2. Homogeneous Volumes

When using distance sampling to find scattering locations in a volume, we are determining not only where each sample might interact, but also whether it does so at all. If we let T_v represent the transmittance through a volume v , and consider a thin volume where $T_v = 0.9$, only 10% of all samples will stop in the volume and thus be used to sample lighting. As volumes tend to converge more slowly than surfaces given identical lighting conditions, this natural distribution of samples may not be suitably relative to the variance in each of the two cases, and surfaces wind up being oversampled long before volumes converge. This is particularly sub-optimal in the common production use case of rendering elements in separate layers with geometry as black holdouts (also known as matte objects); here, we will traverse the volume N times but only compute $N \cdot (1 - T_v)$ light samples, effectively wasting $N \cdot T_v$ samples that do not contribute to the image. The cost of traversal involves not only access to the acceleration structure, but also a potentially large number of evaluations of the underlying density function.

Given a finite volume that extends a distance t_v with $\mu_t = \hat{\mu}$, we can address this imbalance by changing the classical distance sampling method $\log(1 - \xi)/\hat{\mu}$ to restrict the range of distances such that all samples drawn fall within the known extents of the volume. In the case of a homogeneous volume with known extents, this has the trivial closed-form solution $T_v = \exp(-t_v \mu_t)$. We call the distance t_n that defines this range the *normalization distance*, and we show that it can be chosen arbitrarily for $t_n \in (t_v, \infty)$ without biasing the estimate. In the case where volumes are rendered

as a separate element, it is optimal to set $t_n = t_v$, but when volumes and surfaces are mixed, t_n must be chosen to lie beyond t_v , such that some samples may escape the volume and reach the surfaces. Once $t_n < \infty$, the probability density function (PDF) no longer cancels out, as it does in the delta tracking case, and we derive a new normalization factor below.

In order to derive our sampling method, we begin with the radiative transfer equation [Kajiya and Von Herzen 1984]

$$L(x, \omega) = \int_0^\infty \exp\left(-\int_0^s \mu_t(x_{s'}) ds'\right) \mu_s(x_s) L_s(x_s, \omega) ds,$$

which simplifies in the homogeneous case to

$$L(x, \omega) = \int_0^\infty \exp(-\mu_t s) \mu_s(x_s) L_s(x_s, \omega) ds.$$

Knowing that $\mu_s(x_s)$ is null outside the volume ($t_n > t_v$), we have

$$\forall t_n \geq t_v, L(x, \omega) = \int_0^{t_n} \exp(-\mu_t s) \mu_s(x_s) L_s(x_s, \omega) ds.$$

The corresponding probability density function (see Figure 2(a)) proportional to density is

$$p(s) = \mu_t \exp(-\mu_t s),$$

which, when defined in the reduced interval $[0, t_n)$, becomes (see Figure 2(b))

$$p(s) = \frac{\mu_t}{1 - \exp(-\mu_t t_n)} \exp(-\mu_t s).$$

As we can see, the only change is the new normalization factor $1 - \exp(-\mu_t t_n)$ needed to account for the fact that we are integrating a sub-segment of the ray (Algorithm 1).




Figure 2. (a) The probability density function $p(s)$ normalized between 0 and ∞ with a volume present between 0 and 2; (b) $p(s)$ normalized between 0 and 3 with a volume present between 0 and 2. Note that the curve has been re-normalized to integrate to 1 between its bounds.

Algorithm 1 Pseudocode for normalized distance sampling in a homogeneous medium. Here, ξ is a uniform random number and f_p is the medium's phase function.

```

function NORMALIZEDDISTANCESAMPLING( $x, \omega$ )
   $s \leftarrow \frac{\ln(1-\xi \cdot (1-\exp(-\mu_t t_n)))}{\mu_t}$ 
   $w \leftarrow (1 - \exp(-\mu_t t_n))$ 
   $x \leftarrow x + s \cdot \omega$ 
  if  $s \geq t_v$  then
     $w \leftarrow w \cdot \frac{\exp(-\mu_t t_v)}{\exp(-\mu_t t_v) - \exp(-\mu_t t_n)}$ 
  else
     $w \leftarrow w \cdot \frac{\mu_s}{\mu_t}$ 
     $\omega \leftarrow \text{sample} \propto f_p(\omega)$ 

```

For samples ending outside of the volume, since we don't cover the full distance $[t_v, \infty)$ anymore but only $[t_v, t_n)$, we need to re-normalize their contribution by the ratio of the cumulative density function (CDF)

$$c(t_v) = \int_{t_v}^{\infty} \mu_t \exp(-\mu_t s) ds,$$

$$c(t_v) = \exp(-\mu_t t_v).$$

Finally, the CDF for the reduced integral bounds is

$$c(t_v, t_n) = \int_{t_v}^{t_n} \mu_t \exp(-\mu_t s) ds,$$

$$c(t_v, t_n) = \exp(-\mu_t t_v) - \exp(-\mu_t t_n).$$

3. Piecewise Homogeneous Volumes

Before approaching the more general heterogeneous case, we adapt the previous algorithm to piecewise homogeneous volumes. As most production scenes involve varying-density media, we first show that the distance normalization method can be applied in the piecewise homogeneous case, which is how volumes are integrated in the case of using volume aggregates to provide ray segments. We also show how a change of variables allows us to reinterpret the integration entirely in optical depth space,¹ enabling further optimizations.

We first imagine a piecewise constant volume which is exactly aligned with our volume-aggregate acceleration structure. Each leaf node has a different $\hat{\mu}$, but the

¹Optical depth is the product $\tau = \mu s$, which is the inner term of Beer's law. For random (uncorrelated, exponential) media, traveling a distance $s = 2$ in a medium of density $\mu = 1$ is equivalent to traveling a distance $s = 1$ in a medium of density $\mu = 2$.

extinction is constant throughout each cell, i.e., $\mu_t = \hat{\mu}$. Using the linear relationship that exists between distance s and extinction μ_t in Beer's Law, $T = \exp(-s\mu_t)$, we can perform a variable change to completely remove the dependence on μ_t , thus operating in units of optical depth.

The transmittance T at distance t for a ray going through the scene is defined as

$$T(t) = \exp\left(-\int_0^t \mu_t(x_s) ds\right).$$

Assuming that we are going through N nodes, each with a different μ_t and the ray/node intersection length t_i , such that

$$T(t) = \exp\left(-\sum_{i=1}^N \mu_t(x_i) t_i\right),$$

we can normalize each cell with respect to a unit $\mu'_t = 1$ without changing the resulting transmittance by scaling its length t_i by μ_t :

$$T(t) = \exp\left(-\sum_{i=1}^N t'_i\right) \text{ with } t'_i = \mu_t(x_i) t_i.$$

Then, we have

$$T(t) = \exp(-t'(t)) \text{ with } t'(t) = \sum_{i=1}^N t'_i,$$

and we can adapt the normalized distance-sampling algorithm from the previous section again as we have successfully converted our multi-segment traversal into a virtual single-cell traversal with an equivalent transmittance (Algorithm 2), as illustrated in

Algorithm 2 Pseudocode for normalized distance sampling in optical depth space. Prime variables are in optical depth space.

function NORMALIZEDOPTICALDEPTH SAMPLING(x, ω)

$s' \leftarrow \ln(1 - \xi \cdot (1 - \exp(-t'_n)))$

$w \leftarrow (1 - \exp(-t'_n))$

$s \leftarrow \text{convertToRealDistance}(s')$

$x \leftarrow x + s \cdot \omega$

if $s' \geq t'_v$ **then**

$w \leftarrow w \cdot \frac{\exp(-t'_v)}{\exp(-t'_v) - \exp(-t'_n)}$

else

$w \leftarrow w \cdot \frac{\mu_s}{\mu_t}$

$\omega \leftarrow \text{sample} \propto f_p(\omega)$




Figure 3. Space stretching: each cell is normalized to unit extinction with cell size scaled accordingly to maintain unchanged transmittance.

Figure 3. The only additional step is to un-stretch the normalized distance back to the original when computing the actual scattering point x .

The transformation from t' to t , i.e., from optical depth space to physical distance is straightforward; we accumulate the equivalent physical distance $s = s' \hat{\mu}$ each time we step forward a distance s' in optical depth space, which keeps the two variables in sync.

3.1. Optimizing Random Number Generation

An important side effect of this shift is that we can avoid re-drawing random numbers each time we cross a segment boundary. This is similar to Szirmay-Kalos et al. [Szirmay-Kalos et al. 2011] where they perform sampling in optical depth in their 3D-DDA traversal. Yue et al.'s method [Yue et al. 2010; Yue et al. 2011] proves that delta tracking can be implemented in a piecewise fashion, but depends on a new random number being sampled for each interval. In our case, this is no longer necessary, and we can utilize the residual step length even as segment boundaries are crossed. As the results section shows, this can reduce the number of random number calls by more than 50%.

3.2. Heterogeneous Volumes

Because our normalized distance sampling scheme only alters the way step lengths and sample weights are chosen, we can trivially extend it to handle heterogeneous volumes by incorporating the same null collision choice as used in delta tracking (Algorithm 3).

Algorithm 3 Pseudocode for normalized distance delta tracking where ζ and ξ are uniform random numbers and primed variables are in optical depth space. The value $\mu_n = \hat{\mu} - \mu_t$ is the null collision coefficient.

function NORMALIZEDOPTICALDEPTHTRACKING(x, ω)

```

 $w \leftarrow 1$ 
while true do
   $s' \leftarrow \ln(1 - \xi \cdot (1 - \exp(-t'_n)))$ 
   $w \leftarrow w \cdot (1 - \exp(-t'_n))$ 
   $s \leftarrow \text{convertToRealDistance}(s')$ 
   $x \leftarrow x + s \cdot \omega$ 
  if  $s' \geq t'_v$  then
     $w \leftarrow w \cdot \frac{\exp(-t'_v)}{\exp(-t'_v) - \exp(-t'_n)}$  return
  else
     $P_s \leftarrow \frac{\mu_t}{\bar{\mu}}$ 
    if  $\zeta < P_s$  then
       $w \leftarrow w \cdot \frac{\mu_s}{\bar{\mu} \cdot P_s}$ 
       $\omega \leftarrow \text{sample} \propto f_p(\omega)$  return
    else
       $w \leftarrow w \cdot \frac{\mu_n}{\bar{\mu} \cdot (1 - P_s)}$ 
       $t'_n \leftarrow t'_n - s'$ 
       $t'_v \leftarrow t'_v - s'$ 

```

3.3. Null-collision Probability Bias

While the previous algorithm places the expected number of samples within a homogeneous volume, null-collisions reduce this number by $\mu_n/\hat{\mu} = 1 - \mu_t/\hat{\mu}$, a ratio that depends on how heterogeneous the actual volume is. In order to further control the balance between samples that terminate in the volume against those that pass through, we introduce a bias on the null-collision probability (Algorithm 4). The goal is to find an increased probability p' of samples stopping in thin heterogeneous areas of the volume.

Through experimentation we found that a gamma function

$$p' = p^{1/\gamma}$$

works well, but that an estimate \tilde{T} of each ray's transmittance is required in order to isolate the effect to thin areas. There are several ways to compute \tilde{T} : if the spatial acceleration structure itself is fine-grained enough, the integral of $\hat{\mu}$ can be used. In the case of volume aggregates, the split threshold generally produces a subdivision that is too coarse to provide an accurate \tilde{T} , in which case we use ratio tracking to compute the estimate. While this requires two traversals of each camera ray (we only bias

Algorithm 4 Pseudocode for biased and normalized distance delta tracking where ζ and ξ are uniform random numbers and primed variables are in optical depth space. The value $\mu_n = \hat{\mu} - \mu_t$ is the null collision coefficient.

```

function BIASEDNORMALIZEDOPTICALDEPTHTRACKING( $x, \omega$ )
   $w \leftarrow 1$ 
  while true do
     $s' \leftarrow \ln(1 - \xi \cdot (1 - \exp(-t'_n)))$ 
     $w \leftarrow w \cdot (1 - \exp(-t'_n))$ 
     $x \leftarrow x + \text{convertToRealDistance}(s') \cdot \omega$ 
    if  $s' \geq t'_v$  then
       $w \leftarrow w \cdot \frac{\exp(-t'_v)}{\exp(-t'_v) - \exp(-t'_n)}$  return
    else
       $P_s \leftarrow (\frac{\mu_t}{\bar{\mu}})^{1/\gamma'}$ 
      if  $\zeta < P_s$  then
         $w \leftarrow w \cdot \frac{\mu_s}{\bar{\mu} \cdot P_s}$ 
         $\omega \leftarrow \text{sample} \propto f_p(\omega)$  return
      else
         $w \leftarrow w \cdot \frac{\mu_n}{\bar{\mu} \cdot (1 - P_s)}$ 
         $t'_n \leftarrow t'_n - s'$ 
         $t'_v \leftarrow t'_v - s'$ 

```


probabilities for the first bounce of light), this is not a hindrance as a separate computation of transmittance is already performed for the purpose of producing an alpha channel. In order to keep the performance under control while still staying unbiased, our implementation automatically switches from ratio tracking to delta tracking when the estimate gets under a certain threshold. We use this \tilde{T} to find our final biased probability as

$$p' = p^{1/(\gamma')} \text{ with } \gamma' = 1 + (1 - \tilde{T}) \cdot (\gamma - 1).$$

Section 4 shows our experimental results, and also provides an example of the importance of isolating the biased probability to thin parts of the volume. A value of $\gamma = 2$ was found to work well in all of our tests.

4. Results

The presented methods address two orthogonal problems in tracking-based integrators: improving sampling of thin homogeneous volumes, as well as thin heterogeneous volumes, while incorporating the reduction of random number calls. In order to illustrate the behavior of each of our optimizations, we evaluate them in three different heterogeneous scenarios. In each case, we compare renders using 16 samples



(a) Heterogeneous fog.



(b) Thin clouds.



(c) Thick clouds with atmosphere.

Figure 4. Evaluation examples – ground truth. 16 bounces, 4096 samples per pixel).

per pixel (spp) to a ground-truth image using 4096 spp, in order to show both how render times are affected by the algorithmic changes, and also how variance improves at equal sample count.

In the first example (Figure 4(a)), the average accumulated extinction is very low ($T \approx 0.005$), and the volume aggregate does not subdivide past the root node, effectively letting the integrator perform tracking of the whole scene in a single segment. This example highlights the behavior of both the normalization and the gamma function independently of the spatial-acceleration structure, but it does not exercise the random number generator optimization. We find that changing γ from 1 to 2 brings the root mean square (RMS) error down by 19%, equivalent to 24 spp. Enabling only distance normalization (with $t_n = t_v$) reduces RMS error by 63%, with equivalent variance at 128 spp. Finally, with $\gamma = 2$ combined with normalization reduces RMS error by 71%, with equivalent variance at 192 spp, for a total sample efficiency gain of 12x. As expected, the distance normalization provides the greatest gain due to the

Experiment	γ	Normalization	RMS error	Render time
Heterogeneous fog	1	Off	0.063	13.8s
	2	Off	0.051	13.0s
	1	On	0.023	17.0s
	2	On	0.018	15.4s
Thin clouds	1	Off	0.020	75.3s
	2	Off	0.017	73.6s
	1	On	0.019	82.9s
	2	On	0.016	83.0s
Thick clouds	1	Off	0.141	239.8s
	2	Off	0.141	233.3s
	1	On	0.132	279.7s
	2	On	0.132	244.4s

Table 1. RMS error and render times for various combinations of γ and distance normalization. In all three cases, the most efficient combination is $\gamma = 2$ and normalization enabled.

very low overall extinction. Without it, most samples simply step through the entire volume and provide no opportunity for the probability bias to take effect. Figure 5 shows the visual result of the four combinations.

The thin clouds in Figure 4(b) provide an example where the average integral of $\hat{\mu}$ is high, suggesting $\tilde{T} = 0$, but due to both large- and small-scale heterogeneity, the




Figure 5. Heterogeneous fog example (16 samples per pixel).




Figure 6. Thin cloud example. 16 samples per pixel.

actual T is closer to 1. In this case, the distance normalization method alone is unable to place more samples within the volume, and we primarily rely on the probability bias for variance reduction. As shown in Table 1, we get an improvement in RMS error of only 1.6% when enabling distance normalization, 17.7% with $\gamma = 2$, and 19.2% combined, with an equal variance at 24 spp. Figure 6 shows the visual differences.

We note that the probability-biasing method is able to reduce variance in the images without increasing render time, even though a much larger number of light samples and shadowing calculations are performed. This can be understood intuitively as being a simple redirection of the ray. The cost of traversing a ray towards the light source is comparable in magnitude to the cost we would have incurred by continuing our traversal, as would happen if no interaction point was found.

In the case of the thick clouds in Figure 4(c), the normalization method is able to reduce the variance in pixels with thin atmosphere. Neither the probability bias nor the distance normalization is able to improve the sampling of the thick clouds, however, since in areas where $T \approx 0$ each sample is already more or less guaranteed to scatter rather than pass through. In this case, the distance normalization automatically reduces its effect without performance penalty. As discussed previously, we also reduce the value of γ progressively as T approaches zero. Figure 7 shows that our method handles this situation gracefully, without degrading variance in dense areas. Without this rolloff, dense areas would be sampled prematurely, causing the types of firefly samples seen in Figure 8.




Figure 7. Thick cloud example (16 samples per pixel).

We note that Figure 7 shows a subtle halo around the dense cloud. This is due to the granularity of the volume aggregate. In nodes that fall along the edge of the cloud, $\hat{\mu}$ is high and $\tilde{\mu}$ is nearly zero. In this case, the distance normalization must take the conservative value into account when normalizing the potential longest distance to sample in optical space, and we see a smaller improvement in variance than nearby areas with only thin media. Although this is visible at low sample counts, we note that the sampling is no worse than in the control case, and that adaptive sampling easily resolves the difference automatically.




Figure 8. The effect of the gamma function with and without an estimated transmittance. The exposure has been adjusted by -5 f-stops to better show the fireflies in (b).

Finally, we track the number of calls made to the random number generator in each of the three examples. As the noise bank is thin and the integrator receives only a single integration interval, there are no continuations to leverage, and the number of calls stays the same. A single-sample render of the second example yields 16,784,300 calls to the generator in the control case, and only 6,075,793 when re-using the residual of ξ , a reduction of 63.8%. In the case of the thick cloud with surrounding atmosphere, the number of calls drops from 21,024,652 to 9,700,829, or 53.9%. We highlight that this reduction in calls is directly related to the efficiency of the acceleration structure used to provide ray segments to the integrator: if long, poorly bounded segments are given, there is little benefit as many steps are required per segment, while overly short segments that are likely to be immediately stepped-through show greater benefits than the examples above. The benefit is also relative to the cost of the generator itself; depending on the implementation a typical volume renderer may spend between 1% and 10% in random number generation.

5. Conclusion And Future Work

We presented two methods for providing user control over sampling of thin volumes when using tracking-based integrators. The distance normalization allows for complete control over the ratio of samples taken in the volume versus taken on surfaces, and the additional probability bias provides further control for heterogeneous volumes. We highlight the fact that the methods for the homogeneous and heterogeneous cases are independent and handled simultaneously, and that as $\check{\mu}$ approaches $\hat{\mu}$, the method naturally converges to the homogeneous case.

We also showed how performing integration in a normalized optical depth space can help reduce the number of calls to the random number generator by allowing re-use of residual step lengths when crossing integration interval bounds.

In the future, we would like to investigate adaptive methods for automatically choosing the t_n variable in order to optimally allocate a sample budget between surfaces and volumes.

References

- FONG, J., WRENNINGE, M., KULLA, C., AND HABEL, R. 2017. Production volume rendering: Siggraph 2017 course. In *ACM SIGGRAPH 2017 Courses*, ACM, New York, NY, USA, SIGGRAPH '17, 2:1–2:79. URL: <http://doi.acm.org/10.1145/3084873.3084907>. 52
- KAJIYA, J. T., AND VON HERZEN, B. P. 1984. Ray tracing volume densities. *SIGGRAPH Comput. Graph.* 18, 3 (Jan.), 165–174. URL: <http://doi.acm.org/10.1145/964965.808594>. 53

- LOKOVIC, T., AND VEACH, E. 2000. Deep shadow maps. In *Proceedings of the 27th Annual Conference on Computer Graphics and Interactive Techniques*, ACM Press/Addison-Wesley Publishing Co., New York, NY, USA, SIGGRAPH '00, 385–392. URL: <http://dx.doi.org/10.1145/344779.344958>. 51
- NOVÁK, J., GEORGIEV, I., HANIKA, J., AND JAROSZ, W. 2018. Monte Carlo methods for volumetric light transport simulation. *Computer Graphics Forum (Proceedings of Eurographics - State of the Art Reports)* 37, 2 (May). URL: <https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.13383>. 51
- PERLIN, K., AND HOFFERT, E. M. 1989. Hypertexture. *SIGGRAPH Comput. Graph.* 23, 3 (July), 253–262. URL: <http://doi.acm.org/10.1145/74334.74359>. 50
- SKULLERUD, H. R. 1968. The stochastic computer simulation of ion motion in a gas subjected to a constant electric field. *Journal of Physics D: Applied Physics* 1, 11, 1567–1568. 51
- SZIRMAY-KALOS, L., TÓTH, B., AND MAGDICS, M. 2011. Free path sampling in high resolution inhomogeneous participating media. *Computer Graphics Forum* 30, 1, 85–97. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1111/j.1467-8659.2010.01831.x>. 52, 56
- SZIRMAY-KALOS, L., GEORGIEV, I., MAGDICS, M., MOLNÁR, B., AND LÉGRÁDY, D. 2017. Unbiased light transport estimators for inhomogeneous participating media. *Computer Graphics Forum* 36, 2, 9–19. URL: <https://onlinelibrary.wiley.com/doi/abs/10.1111/cgf.13102>. 52
- WOODCOCK, E., MURPHY, T., HEMMINGS, P., AND LONGWORTH, T. 1965. Techniques used in the GEM code for Monte Carlo neutronics calculations in reactors and other systems of complex geometry. In *Applications of Computing Methods to Reactor Problems*, Argonne National Laboratory; available from the Clearinghouse for Federal Scientific and Technical Information, Springfield, VA. 51
- WRENNINGE, M., AND BIN ZAFAR, N. 2011. Production Volume Rendering 1: Fundamentals. ACM, New York, NY, USA, ACM SIGGRAPH 2011 Courses. 51
- WRENNINGE, M., BIN ZAFAR, N., HARDING, O., GRAHAM, G., TESSENDORF, J., GRANT, V., CLINTON, A., AND BOUTHORS, A. 2011. Production Volume Rendering 2: Systems. ACM, New York, NY, USA, ACM SIGGRAPH 2011 Courses. 51
- WRENNINGE, M. 2012. *Production Volume Rendering: Design and Implementation*. A K Peters/CRC Press, Boca Raton, FL. 51
- YUE, Y., IWASAKI, K., CHEN, B.-Y., DOBASHI, Y., AND NISHITA, T. 2010. Unbiased, adaptive stochastic sampling for rendering inhomogeneous participating media. *ACM Trans. Graph.* 29, 6 (Dec.), 177:1–177:8. URL: <http://doi.acm.org/10.1145/1882261.1866199>. 52, 56
- YUE, Y., IWASAKI, K., CHEN, B.-Y., DOBASHI, Y., AND NISHITA, T. 2011. Toward optimal space partitioning for unbiased, adaptive free path sampling of inhomogeneous participating media. *Computer Graphics Forum* 30, 7, 1911–1919. URL: <http://dx.doi.org/10.1111/j.1467-8659.2011.02049.x>. 52, 56

Author Contact Information

Ryusuke Villemin	Magnus Wrenninge	Julian Fong
Pixar Animation Studios	Pixar Animation Studios	Pixar Animation Studios
1215 45th Street	1215 45th Street	1215 45th Street
Emeryville, CA 94608	Emeryville, CA 94608	Emeryville, CA 94608
rvillemin@pixar.com	magnus@pixar.com	jfong@pixar.com

Villemin, Wrenninge, Fong, Efficient Unbiased Rendering of Thin Participating Media, *Journal of Computer Graphics Techniques (JCGT)*, vol. 7, no. 3, 50–65, 2018
<http://jcgt.org/published/0007/03/03/>

Received: 2018-03-12
Recommended: 2018-05-08
Published: 2018-09-13

Corresponding Editor: Wenzel Jakob
Editor-in-Chief: Marc Olano

© 2018 Villemin, Wrenninge, Fong (the Authors).

The Authors provide this document (the Work) under the Creative Commons CC BY-ND 3.0 license available online at <http://creativecommons.org/licenses/by-nd/3.0/>. The Authors further grant permission for reuse of images and text from the first page of the Work, provided that the reuse is for the purpose of promoting and/or summarizing the Work in scholarly venues and that any reuse is accompanied by a scientific citation to the Work.

