Product Importance Sampling of the Volume Rendering Equation using Virtual Density Segments (Pixar Technical Memo 20-01)

MAGNUS WRENNINGE, Pixar Animation Studios
RYUSUKE VILLEMIN, Pixar Animation Studios

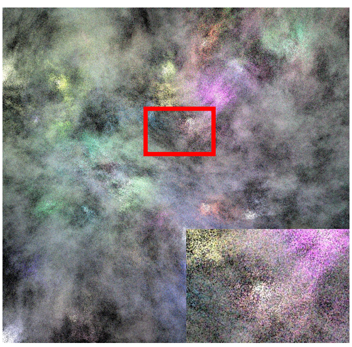

(a) Delta tracking
SMAPE: 0.316

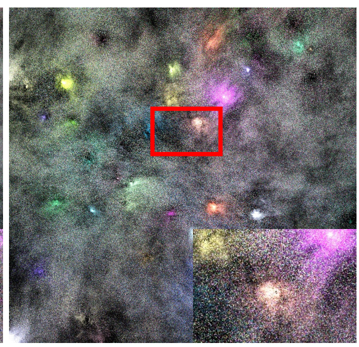

(b) Equiangular sampling
SMAPE: 0.535

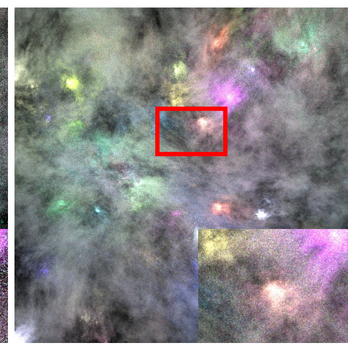

(c) Resampled product importance sampling
SMAPE: 0.185

Fig. 1. Noise bank with 50 colored light sources at fixed render time of 1m per image. Delta tracking can resolve the heterogeneous density (left) and equiangular sampling improves areas around light sources (middle). Our resampled product importance sampling method (right) handles efficient sampling of both the density and light distributions, even in a many-light situation.

1 INTRODUCTION AND PREVIOUS WORK

With path tracing [4] firmly established as the basis for nearly all film-oriented rendering, a large body of research into the related mathematics and integration techniques have followed. While both CPUs and GPUs get faster each year, the need to create ever-more spectacular imagery demands that improvements are made on as many fronts as possible. Given the $\sqrt{N}$ convergence rate of Monte Carlo simulation, variance reduction techniques, such as *importance sampling*, are one of the key paths towards making film-quality image rendering faster.

When rendering participating media such as smoke, clouds, or fire, the volume rendering equation (VRE) defines how light is emitted, scattered, and absorbed within the medium:

$$L(x,\omega) = \int_0^t T(x,y)\Big(\mu_t(y)L_e(y,\omega) + \mu_s(y)p(\omega_i,\omega_o)L_s(y,\omega)\Big)dy. \quad (1)$$

Parts of the VRE can be importance sampled, as has been shown in several works. Tracking methods [7, 12] effectively importance sample the transmittance term; equiangular sampling [5] importance samples the radial falloff of the L_s term, and other works have focused on the emissive L_e term. While it is beneficial to importance sample individual terms compared to sampling blindly, it does not yield sample distributions that conform to the VRE as a whole. To address this problem,

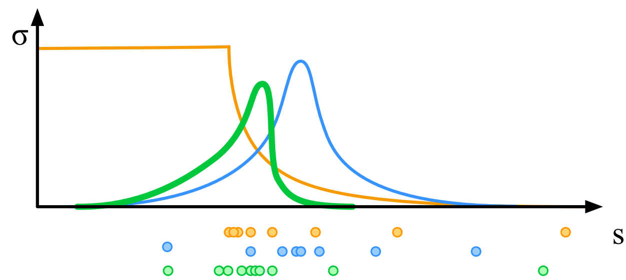

Fig. 2. The function $f(s) = T \cdot L$ in green with T in orange (representing transmittance) and L in blue (representing light contribution). The peak of the product does not line up with either of the two terms. Multiple importance sampling (MIS) can combine samples from multiple distributions by weighting, but does not produce a sample set that conforms to the true product of the distributions.

multiple importance sampling (MIS) [10] provides a way of combining samples drawn from multiple distributions by combining the samples' weights according to both distributions. However, MIS has several shortcomings: the samples produced have weights according to the combined importance, but they are not themselves distributed according to the product of the terms; additionally, MIS requires a way of evaluating the *probability density function* (PDF) for each distribution, something that tracking-based methods such as Woodcock tracking cannot provide. Figure 2 illustrates how MIS can produce samples that miss the true peak of a product importance.

Georgiev et al. [2013] showed a method for importance sampling of multi-vertex paths in anisotropic volumes where sample generation took into account the transmittance along both the camera and shadow rays, as well as the phase function, but was limited to homogeneous media.

Besides not providing a PDF, tracking is inherently a different form of importance sampling than inversion-based techniques. While e.g. equiangular sampling (and many types of BRDF and phase function sampling) rely on inverting the cumulative density function (CDF) mathematically, tracking is iterative and *physically analogous*, reflecting the behavior of individual particles moving through a medium filled with interacting matter. As such, tracking fuses the *mechanics* of the physical problem with the mathematics. When implementing a tracking-based integrator, this analogous behavior of simulating a particle's path along a ray fits well with the way a ray traverses an acceleration structure in the computer's memory, and this is an advantage we want to retain, along with tracking being unbiased, consistent, as well as efficient.

The mechanics of tracking (incrementally taking random steps along a ray and sampling the underlying volume) also explain why we cannot directly compute a PDF value for an arbitrary point on a ray: in order to compute the PDF value, we would need information about the transmittance up to the given point (an iterative process) as well as the transmittance for the entirety of the ray until infinity (continuing the iterative process). With this in mind, it is clear that MIS is ill-suited as a method for importance sampling multiple terms of the VRE using tracking methods.

Another aspect of tracking concerns the sampling majorant $\hat{\mu}$, which controls the length of each tracking step according to

$$s = -\frac{\ln(1 - \xi)}{\hat{\mu}}. \tag{2}$$

When employing a e.g. delta tracking in a direct lighting context, we need to increment the tracking step in this fashion, and at each step determine whether a real or virtual scattering event is found according to

$$P_{real} = \frac{\mu}{\hat{\mu}}. \tag{3}$$

The stepping needs to be continued until a real scattering event is found, so in general we want $\hat{\mu}$ to be a tight bound on μ, such that the P_{real} is large. In fact, many important works explore methods for optimizing this aspect of tracking [1, 13], and a good acceleration structure is key to making tracking practical in production rendering.

2 INTEGRAL FORMULATION OF TRACKING

Recent work [2, 7] has shown that the notion of using virtual particles to facilitate tracking has a corresponding integral formulation, which also generalizes the probability definitions to arbitrary weights. The integral formulation is helpful in understanding how importance resampling fits with the tracking framework, and that resampling is a natural incorporation of more information into the generalized integral.

As we recall, delta tracking adds fictitious particles into the medium which exactly fills in the difference between the true density μ and the majorant $\hat{\mu}$. We can incorporate this into the VRE by defining the fictitious (virtual, or null) particles as being present in the medium, but scattering light exactly in the direction of travel:

$$\mu_n(x)L(x,\omega) = \mu_n(x)\int_{S^2} \delta(\omega - \hat{\omega})L(x,\hat{\omega})d\hat{\omega}. \tag{4}$$

This leads to a modified volume rendering equation

$$L(x,\omega) = \int_0^\infty T_{\hat{\mu}}(x,y)\Big(\mu_a(y)L_e(y,\omega) + \mu_s(y)L_s(y,\omega) + \mu_n(y)L(y,\omega)\Big)dy, \tag{5}$$

where $T_{\hat{\mu}}(x,y) = e^{-\int_0^y \hat{\mu}(x,z)dz}$ and $\hat{\mu}(z) = \mu_a(z) + \mu_s(z) + \mu_n(z)$.

From this definition, we can then define a Monte Carlo estimator of the equation on the form

$$\langle L(x,\omega)\rangle = \frac{T_{\hat{\mu}}(x,y)}{p_{\hat{\mu}}(y)}\{\mu_a(y)L_e(y,\omega) + \mu_s(y)L_s(y,\omega) + \mu_n(y)L(y,\omega)\}, \tag{6}$$

where we stochastically choose one of the interaction types based on

$$P_a(y) = \frac{\mu_a(y)}{\hat{\mu}(y)}, P_s(y) = \frac{\mu_s(y)}{\hat{\mu}(y)}, P_n(y) = \frac{\mu_n(y)}{\hat{\mu}(y)}. \tag{7}$$

3 TOWARD PRODUCT SAMPLING

The iterative process of finding scattering events using delta tracking produces a sample distribution that is proportional against the $T(x,y)\cdot\mu_t(y)$ terms in the VRE. $T(x,y)$ is not directly evaluated, but rather it is the implicit product of the previous series of virtual scattering probabilities

$$P_{virtual} = 1 - P_{real} = 1 - \frac{\mu}{\hat{\mu}}, \tag{8}$$

against which delta tracking is effectively doing a russian roulette termination using Equation 3. Thus, after i steps along a ray, the estimated transmittance up to that point can be found through

$$T_i = \prod_{j=0}^{i-1} P^j_{virtual}. \tag{9}$$

It may be tempting to try and introduce additional terms from the VRE, such as the radial falloff of the L_s term, thereby incorporating equiangular sampling, but because the termination probabilities

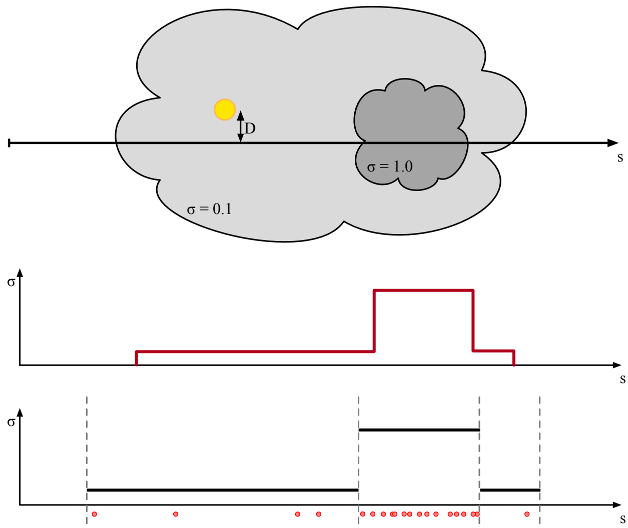

Fig. 3. Top: Example scene with heterogeneous medium and embedded light source. Middle: density profile along camera ray. Bottom: integration segments with $\hat{\mu}$ majorant (bold) and tracking-based sample points (red).

indirectly drive our estimate of T, this is not possible. However, as a thought experiment, we can see that the samples produced by delta tracking are the same as if we had taken a more circuitous path: we could hypothetically perform tracking along the entirety of the ray, and for each step i record a weight w_i as the product of T_i from Equation 9 and P_{real}. The set of locations and their corresponding weights would then define a discrete PDF from which we could randomly draw samples. This process would be more costly than delta tracking, and would yield no better sampling; however it does offer a place for us to incorporate other sampling weights into the decision of where to sample direct lighting: by multiplying each w_i by $\mu_s L_s^i$, we could pick samples according to the full product of the VRE.

In the Monte Carlo literature, the process of computing sample weights by evaluating a function $f(x)$ directly, rather than relying on its PDF, is referred to as *sampling-importance resampling* (SIR), or just *importance resampling* [8]. In graphics, the method has been used for direct lighting on surfaces [9] and shows an alternative way to MIS for combining samples from multiple distributions.

While the importance resampling method of incorporating L_s into the sample weights appears promising at first glance, it suffers from the fact that the *candidate points* that the incremental stepping produces is proportional to $\hat{\mu}$, but not to L_s. Thus, for a ray that passes near a light source, we may never sample L_s at its most important location. Figure 3 shows this arrangement: as the light is embedded in a thin part of the volume, very few sample points fall in the area where L/r^2 is large. In order for importance resampling to be viable in the tracking context, this deficiency would have to be addressed.

Finally, we note that when tracking, the distribution of *real scattering events* generated conforms to the full $T(x, y) \cdot \mu_s(y)$ function of the heterogeneous medium. However, the distribution of *candidate points*, i.e. the points generated by repeated steps through the volume in optical depth space is by definition a uniform, i.i.d. distribution[1], an important factor that we take advantage of in this work. This is again illustrated in Figure 3, where the concentration of candidate points is highest in the dense part of the volume, but within each constant-$\hat{\mu}$ interval, the set of candidate points is uniform. The method presented by Villemin et al. [2018] of *stretching* piece-wise constant segments during tracking such that $\hat{\mu} = 1$ shows that we can treat the candidate points for an arbitrary set of integration intervals as an i.i.d. distribution of samples in *optical depth space*. As such, it is clear that $\hat{\mu}$ acts as a control on the placement of candidate points, and that we can steer an arbitrary density of candidate points to a given segment of a ray simply by increasing $\hat{\mu}$ in that region.

4 OVERVIEW

In order to overcome the previously mentioned constraints on importance sampling, we propose a new volumetric integration method that combines guiding of candidate point positions and importance resampling. We refer to this as the *virtual density segment* method (VDS). The key insight that enables our method is that tracking-based integration can be indirectly controlled using $\hat{\mu}$: in areas where we want a greater concentration of candidate sample points, we simply increase $\hat{\mu}$, and vice versa. In particular, we show that this control can be driven by treating invertible PDFs as *virtual density* sources, which in turn steers a tracking algorithm to generate distributions of points that conform to the same, arbitrary PDFs. We combine this virtual density process with *importance resampling* to pick samples from the set of candidates according to the full product of the VRE. The resampling step is especially beneficial for non-invertible terms, such as complex light shapers like projected texture maps. In the end, by bridging tracking methods and inversion-based importance sampling, we arrive at a method for steering sampling that can incorporate any number of PDFs, thereby providing a general framework for combining arbitrary importance sampling schemes with tracking.

Finally, having employed the importance resampling method for direct lighting, we also introduce a related method to the sampling of indirect volumetric illumination for highly anisotropic media.

5 PROBABILITY DENSITY AS CONTROL SIGNAL FOR TRACKING

Next, we explore in more detail how $\hat{\mu}$ can be used to control sampling distributions. If we consider equiangular sampling, we have the PDF

$$P(t) = \frac{D}{(\theta_b - \theta_a)(D^2 + t^2)} \tag{10}$$

and the corresponding CDF inverse which is used for drawing samples from the distribution:

$$C^{-1}(\xi) = D \tan((1 - \xi)\theta_b + \xi\theta_a). \tag{11}$$

The square falloff function itself, along with some sample points from C^{-1} can be visualized as:

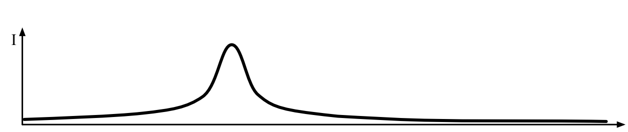

[1] This may seem counter intuitive at first, but although the *distance between points* produced by a tracking algorithm is an exponential distribution, the *set of points* formed by the incremental stepping are in fact a uniform, i.i.d. distribution. See "Memoryless property" at https://en.wikipedia.org/wiki/Poisson_point_process

The domain of C^{-1} is $[0, 1]$, and the range is defined by the subtended angles θ_a and θ_b. In order to produce integration intervals for our tracking algorithm, we divide the domain of C^{-1} into N *equal-size* segments defined as $I_i = [p_i, p_{i+1}]$, with $p_i = \frac{i}{(N-1)}$, which gives us a set of N *equal-importance* segments along a ray.

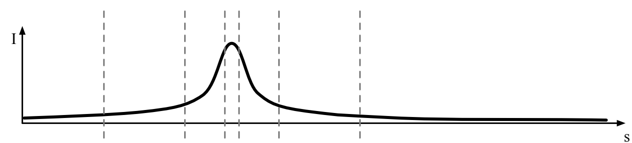

Next, we need to set a $\hat{\mu}$ for each segment such that the tracking algorithm places the desired amount of samples in each one. We know that tracking takes a number of steps through a segment proportionally to $\mu \cdot l$, from which we conclude that we can get an equal likelihood of sampling a given segment if we assign an even *optical thickness* to each segment. Thus, we define

$$\hat{\mu}_i = \frac{c}{l_i}, \tag{12}$$

where c is an arbitrary scaling constant, and l_i is the physical-space length of each segment. Throughout this work, we typically set $c = 1$, which makes each control segment exactly one mean free path in optical thickness. We refer to these control densities as *virtual density segments*.

Given these segments, we are able to achieve a distribution of candidate points in our tracking method that approximates the distribution of an arbitrary PDF: in this case the equiangular distribution, as illustrated in Figure 4.

Although N would need to be large for the resulting distribution to closely resemble the discretized PDF, in practice, the number of segments does not need to be very large. Because the subdivision is performed in the $[0, 1]$ domain of the importance function itself, the resulting intervals automatically adapt to the appropriate areas of interest. If we consider a set S consisting of sample points drawn from C^{-1} and denote the set of points generated by our method S', then the distribution of S' will converge on S as N increases. In fact, we can see that our method is analogous to a Riemann sum, and we are thus guaranteed to produce an S' distribution that converges to the original S distribution.

5.1 Combining control segments

Because our control method works by specifying density intervals, we can incorporate it with the general tracking strategy seen in Figure 3. Given two sets of overlapping intervals, A and B, we construct a new set C by taking the maximum of $\hat{\mu}_a$ and $\hat{\mu}_b$. Figure 5 illustrates the resulting interval set, which is readily employed by the tracking part of our product sampling method.

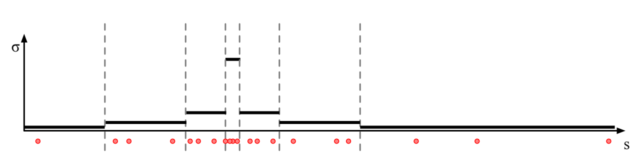

Fig. 4. Seven intervals generated by inversion sampling of the equiangular distribution yield control segments which force tracking methods to generate sample distributions that are approximately proportional to the original distribution.

The motivation for using max rather than addition or multiplication is that the parameter N that defines the number of subdivisions of the control PDF acts as a control over how many times we would like to evaluate the PDF (if $c = 1$ is used, we will get on average $N/2$ samples.) If we were to add the virtual densities to the natural $\hat{\mu}$, certain regions would be sampled more finely than either of the two distributions would suggest. Multiplying the two could cause undersampling of the density portion if the virtual density is below one, which would introduce large amounts of variance into the integration. By taking the maximum of the two densities, we ensure that the tracking process samples *at least* as finely as either of the two distributions desire.

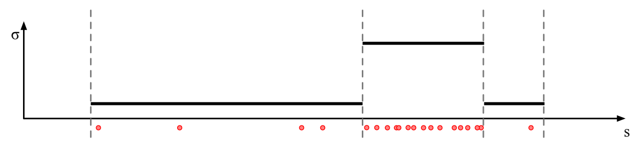

(a) Integration intervals A based on volumetric density.

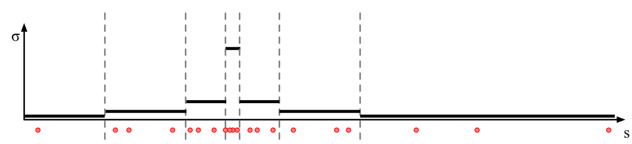

(b) Integration intervals B based on $1/r^2$ falloff.

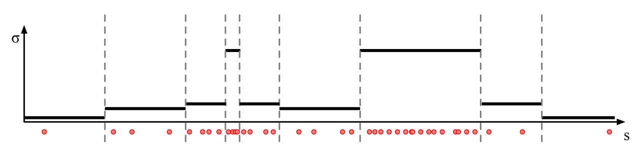

(c) Combination of integration intervals using $\max(A, B)$.

Fig. 5. Multiple sets of integration intervals can be combined in order to capture multiple terms in the VRE.

6 IMPORTANCE RESAMPLING

With a method in place for generating the necessary candidate sample locations, we next show how importance resampling can be used in order to select the most suitable location at which to compute direct illumination. As previously mentioned, we wish to incorporate transmittance, illumination estimate, local scattering properties as well as the phase function response in this decision. To this end, we assign the following weight to each sample point:

$$W_i = T_i \cdot \mu_s(y_i) p(\omega) L_s(y_i, \omega), \tag{13}$$

where T_i is the transmittance up to sample i (Equation 9), μ_s is the scattering coefficient, $p(\omega)$ is the phase function response relative to the light's direction ω and L_s is the I/r^2 intensity of the picked light at the point y_i. We note that none of these require any global information, i.e. they can be sampled directly or computed inexpensively.

Given the set of samples with weights $W = (W_0, W_1, ..., W_i)$, we compute a discrete CDF from which we draw the final sample stochastically.

7 LIGHT PICKING

So far we have only considered scenes with a single light sources, but production scenes often contain hundreds or thousands of light sources. Whether surface or volume rendering is considered, it is computationally impractical to sample the light contribution of all lights. Instead, a process known as *light picking* aims to reduce this two a single light sample for each direct lighting calculation. Light picking is an area of current research [**?**], but for the purposes of this paper it is sufficient to assume that a method for intelligently picking a single relevant light source given a ray is available. In this paper, we use a simple method by which each light is given a contribution likelihood according to

$$C(i, y, \omega) = \frac{I}{\text{DistanceToRay}(p_i, y, \omega)^2}, \tag{14}$$

from which a single light source is chosen by constructing a discrete CDF. The chosen light is then used for computing the virtual density segments along the ray, as well as for estimating the illumination term in Equation 13. However, we note that a new light sample is chosen when computing direct illumination, as point-based light picking can be done more accurately than ray-based picking.

While the one-light approach to constructing the virtual density segments captures the sharp transition in illumination near a light source, the sum of all illumination at a given point is generally much more gradual. To this end, we make the sample weight L_s in Equation 13 less aggressive by taking the logarithm of the illumination estimate according to

$$L_s(y_i, \omega) = \log\left(1 + \frac{I}{r^2}\right). \tag{15}$$

7.1 Distant light sources

We also note that lights far away from scene (such as a sun light or a sky dome light) present a degenerate case for equiangular sampling. In fact, the contribution of these lights (not considering visibility) is the same for the length of the entire ray. Although these infinite-distance lights appear problematic, we note that they are the limit case of the equiangular distribution for very large values of D. As such, in the cases where a light sample is chosen on an infinitely-far light source, we replace the virtual density segment generation step in Equation 11 with a uniform subdivision of the integration interval $[t_0, t_1]$. Besides avoiding the edge case for equiangular sampling, this approach has the additional benefit of ensuring that (on average) at least $N/2$ tracking steps are taken along the ray. This implicitly addresses the common problem with delta tracking in thin volumes, where $\hat{\mu}$ would otherwise be low enough that few samples are computed in the volume. This problem was partially addressed in Villemin et al. [2018], but our method provides a better solution for heterogeneous volumes.

8 OPTIMIZATIONS

We note that our method, in its naïve form, can be inefficient in certain types of scenes. Where densities are large enough to cause transmittance along a ray to approach zero, we will generate many candidate samples for which the weight is also close to zero. This is wasteful, as each sample deep inside of a volume is as expensive as the first but will not contribute significantly.

A common technique in Monte Carlo integration is to use Russian roulette to terminate event sequences that have low contribution to the final estimate. We adopt an analogous process in our construction of the discrete CDF. First, we ensure that the I/r^2 term is monotonically decreasing, i.e. we have passed the point along the ray that is nearest to the light sample. From there, we set a termination threshold K_t (usually 0.01) and randomly terminate the sampling of μ if T_i if it falls below K_t according to

$$\xi < \frac{\mu}{\hat{\mu}} K_t. \tag{16}$$

While this introduces variance (as Russian roulette always does), it can still be worthwhile, as dense regions can require very large numbers of steps to traverse completely.

9 IMPORTANCE RESAMPLING FOR INDIRECT LIGHTING

In production rendering, the term that specifies incoming light in both the volume rendering and surface rendering equations is generally split into two components: one for light arriving directly (without scattering) from a light source, and one arriving indirectly (having scattered at least once) from a light source. This process is called *next event estimation*, and is premised on the fact that much, if not most, of the light arriving at a given point in a scene arrives unimpeded from a light source.

The method described so far in this paper is concerned with deciding *where* along a ray, passing through a participating medium, this next event estimation process should take place. That is to say, it steers samples towards areas that are expected to provide good throughput for direct illumination. However, importance resampling offers a way to improve upon the second part as well: the indirect illumination.

When splitting incoming illumination in the VRE into direct and indirect illumination, we choose the direction for direct illumination by sampling a light. The indirect direction is chosen by randomly generating a direction from the phase function of the volume. We note that the phase function concerns itself only with the distribution of the forward/backward angle compared to the preceding direction of travel, but that the function is entirely symmetric *around* the direction of travel. Again, this occurs because we do not factor any knowledge about the direction of incoming indirect illumination into the equation. Although it is difficult to incorporate the sum of all light sources as a weight on the direction to choose, we note that the direction of direct illumination is likely to serve as a good guess for the general direction from which indirect illumination would arrive from.

In this context we apply a second importance resampling approach: by generating N *candidate directions* d_i from the phase function inversion, we produce a uniform distribution of directions in the plane that is tangent to the direction of travel. Next, we compute a weight w_i for each d_i direction according to the phase function response between d_i and the direction of the previously picked light source, ω_i according to

$$w_i = p(\omega_i, \dot{i}). \tag{17}$$

Similarly to the product sampling technique for direct illumination above, we build a discrete CDF from the weights and pick a single indirect direction, which is drawn directly from the phase function, but which also incorporates knowledge about which directions are likely to produce the greatest indirect illumination contribution.

10 RESULTS

In this section we look at numerical results comparing our two proposed methods compared to existing approaches. Throughout the comparison we have chosen to use Symmetric Mean Absolute Percent Error (SMAPE) rather than Root Mean Square Error (RMSE), which is generally more common for Monte Carlo-based methods. However, RMSE measures absolute error, which means that an over-estimate of 10x (e.g. 100.0 instead of the true value 10.0) produces a much greater error than an under-estimate of 10x (1.0 instead of 10.0). Figure 1 provides a good example of this: the type of noise that is introduced by both delta tracking and equiangular sampling tends to be of under-estimation. Delta tracking can miss the bright areas around a visible light source, but its error is one of omission, for which the error is bounded to the magnitude of the ground-truth value. Likewise, equiangular sampling can produce areas where nearly all chosen samples are so deep inside a volume that they produce zero throughput to the viewer. In both cases, our method may produce an RMSE value that is only slightly lower than the two other methods, even though the visual result is significantly cleaner.

A second reason for using SMAPE rather than RMSE is that it makes all pixels count evenly. When using RMSE, the majority of the value will be contributed by the bright pixels in the image, which under-represents the visual noise in the midtones while over-representing the noise in very bright areas, which (due to the logarithmic nature of human vision) is less noticeable, relatively, than the midtones.

10.1 Virtual density segment sampling

In Table 1 we show the performance of our proposed method (VDS) compared to delta tracking and equiangular sampling in a variety of scenarios. All tests were performed by running the algorithms on an 8-core Xeon workstation for 2 minutes and computing the error on this equal-time basis.

The first table shows results for a single light source, effectively ruling out variation due to the discussed light picking method. The second table show the same scene with 50 light sources. In both cases the results are consistent. In order to show that our method works independently of the light picking method, we also test strongly forward scattering and backward scattering media ($g = 0.95$ and $g = -0.95$, respectively) in addition to an isotropic medium. We also vary the mean free path (MFP) from 0.2, to 1.0 and 10.0, corresponding to scattering coefficients of 50.0, 1.0 and 0.1. Across all the tests, our method outperforms both of the previous methods, with far lower variance in nearly all cases. Figure 6 shows visual examples from the test, illustrating the improvement in perceived variance.

Certain combinations illustrate the weaknesses of the previous methods. Due to the fact that a large number of rays traverse the volume without finding a scattering event, delta tracking performs worst in thin media (the MFP=10.0 case), even though the number of samples taken per pixel is very high (around 800spp, as compared to 150spp for equiangular sampling and 100spp for VDS.) Equiangular sampling fares worst in anisotropic media, since the sample locations chosen near the light source no longer corresponds to high-throughput paths.

10.2 Indirect sampling

Table 2 shows the performance of our proposed method compared to naïve sampling. In the experiment, four candidate directions were generated at each indirect scattering event, which has a negligible impact on performance. As expected, the two methods produce the same amount of variance for $g = 0$, as isotropic scattering give all generated directions the same weight. As g increases, we note that the SMAPE value increases much more slowly for our importance resampling method than the naïve sampling. Figure 7 shows the visual results of the two methods. In our experience, anisotropy values over 0.875 are impractical to render in production scenes with the standard method, as fireflies become too prominent. With our method, it is possible to push g as high as 0.95 while still producing relatively well-behaved variance characteristics.

11 DISCUSSION

As our results illustrate, our proposed method for sampling direct lighting works across a wide range of participating media contexts. Thin and thick media are both handled well, as is strongly anisotropic media and many-lights situations. However, there are some cases in which our product

(a) Single light source

g	MFP	Method	SMAPE
-0.95	0.2	Delta Tracking	0.430
		Equiangular	0.941
		VDS	**0.279**
	1.0	Delta Tracking	0.328
		Equiangular	0.384
		VDS	**0.152**
	10.0	Delta Tracking	0.420
		Equiangular	0.201
		VDS	**0.137**
0	0.2	Delta Tracking	0.347
		Equiangular	0.795
		VDS	**0.224**
	1.0	Delta Tracking	0.285
		Equiangular	0.270
		VDS	**0.120**
	10.0	Delta Tracking	0.373
		Equiangular	0.142
		VDS	**0.092**
0.95	0.2	Delta Tracking	0.382
		Equiangular	0.909
		VDS	**0.250**
	1.0	Delta Tracking	0.322
		Equiangular	0.353
		VDS	**0.151**
	10.0	Delta Tracking	0.430
		Equiangular	0.192
		VDS	**0.128**

(b) 50 light sources

g	MFP	Method	SMAPE
-0.95	0.2	Delta Tracking	0.738
		Equiangular	1.260
		VDS	**0.624**
	1.0	Delta Tracking	0.734
		Equiangular	0.929
		VDS	**0.513**
	10.0	Delta Tracking	0.862
		Equiangular	0.784
		VDS	**0.567**
0	0.2	Delta Tracking	0.326
		Equiangular	0.955
		VDS	**0.264**
	1.0	Delta Tracking	0.316
		Equiangular	0.535
		VDS	**0.185**
	10.0	Delta Tracking	0.422
		Equiangular	0.409
		VDS	**0.216**
0.95	0.2	Delta Tracking	0.825
		Equiangular	1.320
		VDS	**0.709**
	1.0	Delta Tracking	0.833
		Equiangular	0.978
		VDS	**0.605**
	10.0	Delta Tracking	0.954
		Equiangular	0.834
		VDS	**0.623**

Table 1. Results for delta tracking, equiangular sampling and VDS at fixed render time of 2m per image. Error calculated using symmetric mean absolute percentage error (SMAPE).

g	Method	SMAPE
0	Naive sampling	0.328
	Importance resampling	0.328
0.5	Naive sampling	0.339
	Importance resampling	**0.331**
0.75	Naive sampling	0.372
	Importance resampling	**0.351**
0.875	Naive sampling	0.417
	Importance resampling	**0.380**
0.935	Naive sampling	0.455
	Importance resampling	**0.414**

Table 2. Results for naïve and importance-resampled indirect illumination at fixed render time of 2m per image. Error calculated using symmetric mean absolute percentage error (SMAPE).

sampling approach is unable to sample better than basic delta tracking: for a thick medium that is lit only by distant lights, delta tracking and product sampling produce the same distribution of scattering locations. In this case, product sampling only adds to the overhead of selecting a point at which to scatter, without providing any variance reduction. In this case it is still preferable to

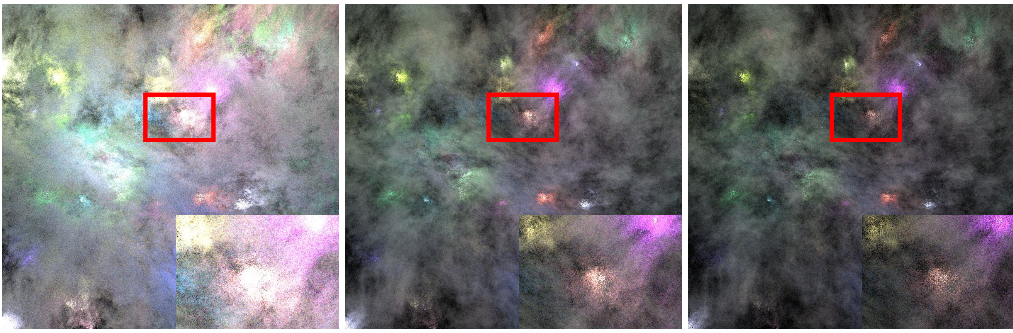

(a) Delta tracking

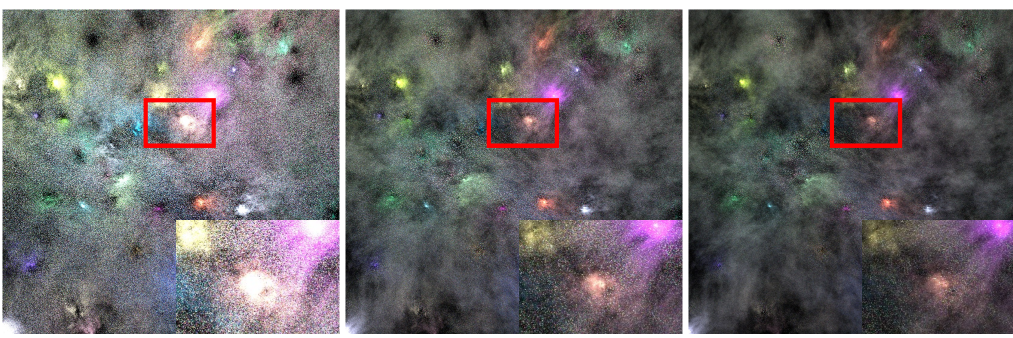

(b) Equiangular sampling

(c) VDS

Fig. 6. Images corresponding to results in Table 1. Left to right: Mean-free-path 0.2, 1.0 and 10.0.

use delta tracking. However - for thick volumes lit by light sources near the volume, our method is still advantageous, as is the case for thin volumes lit by distant lights, since our method is able to find candidate scattering locations even when delta tracking may advance through the whole volume in a single step.

As part of our exploration, we attempted to incorporate the visibility term for light sources into the resampling weight W_i by traversing the volume aggregate and using $\exp(-\int_0^s \hat{\sigma} dt)$ as an estimate of $\exp(-\int_0^s \sigma dt)$ for shadow rays. This proved to increase variance significantly, as

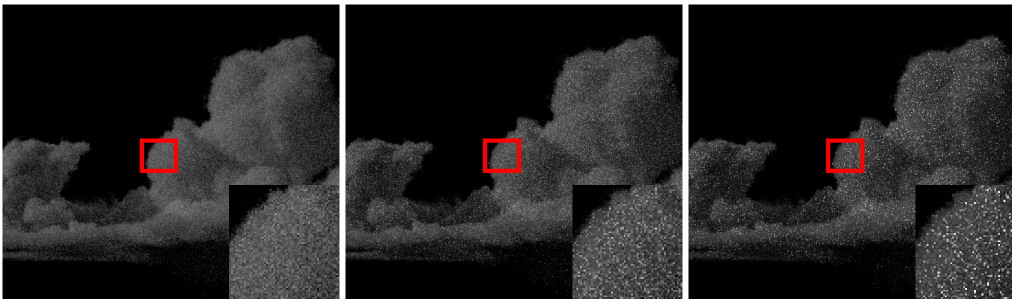

(a) Naive sampling

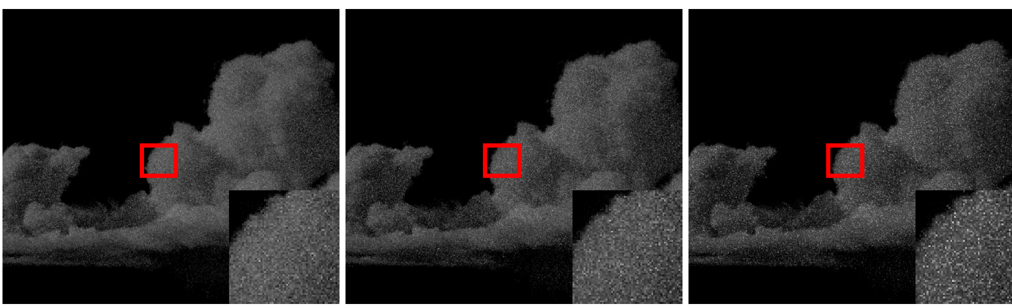

(b) Importance resampling

Fig. 7. Left to right: Anisotropy 0.75, 0.875 and 0.925.

visibility was consistently under-estimated and samples that returned a higher than estimated visibility wound up having both a low PDF value and high throughput.

12 CONCLUSIONS AND FUTURE WORK

We have presented a novel method for importance sampling of direct lighting in participating media that unifies tracking with other, arbitrary sampling PDFs through the use of virtual density segments and importance resampling. The method does not depend on any novel data structures and can be incorporated into existing volume renderers that provides the ability to perform ray-based queries for light picking.

Since the proposed method is a framework that allows arbitrary PDFs to be incorporated, there are several promising areas for future work. One practical example would be to guide sampling according to the intensity of a projected texture (commonly referred to as a *gobo*), which can create artistically pleasing structure in light sources, but which is traditionally impractical to sample. In the virtual density segment case, we can take advantage of the fact that a linear ray in an orthogonal space preserves its linearity when undergoing perspective projection, and we find a 2D line in texture space corresponding to the world-space ray. For each ray, we could perform a 2D DDA walk through the texture in order to create a PDF representing its intensity, which is then integrated and inverted in order to find regions that, when sampled uniformly in importance-space, yields the relevant segments along the ray to focus samples upon. Similarly, volumetric emission could be used to drive importance by building a hierarchical data structure for querying the majorant $\hat{L_e}$ over a given domain, which could be used to construct a piece-wise constant PDF. Finally, we believe that recent work on learning-based methods [6] for importance sampling could be incorporated as a replacement for the full $\frac{I}{r2} \cdot V$ term in the importance resampling weight.

REFERENCES

[1] Julian Fong, Magnus Wrenninge, Christopher Kulla, and Ralf Habel. 2017. Production Volume Rendering: SIGGRAPH 2017 Course. In *ACM SIGGRAPH 2017 Courses (SIGGRAPH '17)*. ACM, New York, NY, USA, Article 2, 79 pages. DOI: http://dx.doi.org/10.1145/3084873.3084907

[2] M. Galtier, S. Blanco, C. Caliot, C. Coustet, J. Dauchet, M. El Hafi, V. Eymet, R. Fournier, J. Gautrais, A. Khuong, B. Piaud, and G. Terrée. 2013. Integral formulation of null-collision Monte Carlo algorithms. 125 (Aug. 2013), 57–68. DOI: http://dx.doi.org/10.1016/j.jqsrt.2013.04.001

[3] Iliyan Georgiev, Jaroslav Kivnek, Toshiya Hachisuka, Derek Nowrouzezahrai, and Wojciech Jarosz. 2013. Joint Importance Sampling of Low-Order Volumetric Scattering. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 32, 6 (nov 2013). DOI: http://dx.doi.org/10.1145/2508363.2508411

[4] James T. Kajiya. 1986. The rendering equation. In *Proceedings of the 13th annual conference on Computer graphics and interactive techniques (SIGGRAPH '86)*. ACM, New York, NY, USA, 143–150. DOI: http://dx.doi.org/10.1145/15922.15902

[5] Christopher Kulla and Marcos Fajardo. 2012. Importance Sampling Techniques for Path Tracing in Participating Media. 31, 4 (June 2012), 1519–1528.

[6] Thomas Müller. 2019. "Practical Path Guiding" in Production. In *ACM SIGGRAPH Courses: Path Guiding in Production, Chapter 10*. ACM, New York, NY, USA, 18:35–18:48. DOI: http://dx.doi.org/10.1145/3305366.3328091

[7] Jan Novák, Iliyan Georgiev, Johannes Hanika, and Wojciech Jarosz. 2018. Monte Carlo Methods for Volumetric Light Transport Simulation. *Computer Graphics Forum (Proceedings of Eurographics - State of the Art Reports)* 37, 2 (may 2018).

[8] Donald B. Rubin. 1987. The Calculation of Posterior Distributions by Data Augmentation: Comment: A Noniterative Sampling/Importance Resampling Alternative to the Data Augmentation Algorithm for Creating a Few Imputations When Fractions of Missing Information Are Modest: The SIR Algorithm. *J. Amer. Statist. Assoc.* 82, 398 (1987), 543–546. http://www.jstor.org/stable/2289460

[9] Justin F Talbot. 2005. Importance resampling for global illumination. (2005).

[10] Eric Veach. 1997. *Robust Monte Carlo methods for light transport simulation.* Vol. 1610. Stanford University PhD thesis.

[11] Ryusuke Villemin, Magnus Wrenninge, and Julian Fong. 2018. Efficient Unbiased Rendering of Thin Participating Media. *Journal of Computer Graphics Techniques (JCGT)* 7, 3 (13 September 2018), 50–65. http://jcgt.org/published/0007/03/03/

[12] E.R. Woodcock, T. Murphy, P.J. Hemmings, and T.C. Longworth. 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.

[13] Yonghao Yue, Kei Iwasaki, Bing-Yu Chen, Yoshinori Dobashi, and Tomoyuki Nishita. 2010. Unbiased, Adaptive Stochastic Sampling for Rendering Inhomogeneous Participating Media. *ACM Trans. Graph.* 29, 6, Article 177 (Dec. 2010), 8 pages. DOI: http://dx.doi.org/10.1145/1882261.1866199

AUTHOR CONTACT INFORMATION

Magnus Wrenninge
magnus@pixar.com

Ryusuke Villemin
rvillemin@pixar.com