MATHEMATICA MONTISNIGRI
Vol XLIV (2019)

CALCULATION OF LUMINANCE OF SCATTERING MEDIUM BY MCRT USING MULTIPLE INTEGRATION SPHERES

S. V. ERSHOV[1], D. D. ZHDANOV[2], A. G. VOLOBOY[1]

[1]Keldysh Institute of Applied Mathematics of RAS,
Miusskaya Sq. 4, Moscow, Russia, 125047
e-mail: voloboy@gin.keldysh.ru, web page: http://keldysh.ru/

[2]ITMO University, 49 Kronverksky Pr., St. Petersburg, 197101, Russia

DOI: 10.20948/mathmon-2019-44-10

Summary. We propose and investigate calculation of the radiance of scattering medium using an extension of bi-directional Monte-Carlo ray tracing with photon maps. In the standard approach photons are collected either by an integration sphere at the end of camera ray segment, or by a cylinder along it. We propose an extension of random ray generation which sets *several* integration spheres per segment of camera ray. Their positions and number are random. The control parameter is their mean number per segment; when it is 0, we have the standard method with single sphere while when it gets very large we effectively obtain the integration cylinder. Although in case of multiple spheres the noise for the *same number* of traced rays is higher than for integration cylinders, processing of forward rays is not that expensive, so for the *same time* of ray tracing the multiple spheres are advantageous. The control parameter allows to achieve minimal noise by finding the compromise between the speed of ray tracing and the noise per given number of rays.

1 INTRODUCTION

A powerful method of calculation of a virtual camera image is a bi-directional Monte-Carlo ray tracing. Its basic idea is to operate trajectories that always connect light source and camera. They are constructed by "concatenating" two halves: one is from the camera and is obtained by the backward Monte-Carlo ray tracing (BMCRT) and another is from light source and is obtained by the forward Monte-Carlo ray tracing (FMCRT). Then these pieces are "merged" somehow which creates trajectory that connects light source and camera. There are two basic approaches for merging these halves and calculate the visible scene luminance.

The first approach supposes the unbiased bidirectional ray tracing method [1, 2], which does not introduce an additional error of luminance estimation in the calculation of illumination at the observation point. This approach is rather expensive to implement

2010 Mathematics Subject Classification: 78-04, 65C05, 65C20.
Key words and Phrases: scattering medium, bi-directional ray tracing, Monte-Carlo ray tracing.

and can add a significant noise in the calculation of luminance of the caustic component of illumination.

The second method is based on the of photon maps and introduces some bias into the luminance estimation caused by averaging of the illumination inside the integration sphere about the observation point. There are several variants of the method, one being the photon map visualization [3, 4, 5, 6]. Most of these methods use a "classical" approach based on the calculation of global illumination in the form of photon maps, which is visualized as the luminance of secondary and caustic illumination [3, 5, 7].

An alternative approach implements the reverse calculation scheme, namely, generates a visibility map as spheres of the illuminance integration in the direction of observation, which are "filled" with the light photons related to the caustic and the secondary illumination [4, 6]. Here the camera ray is traced stochastically until it terminates due to some criterion, say, after the given number of diffuse events. In each point of scattering, an "integration sphere" is set that collects FMCRT rays. After a forward ray hits that sphere we calculate the surface luminance for the view direction equal to the camera ray direction before the sphere center. This luminance is then scaled by the camera ray attenuation accumulated to this point and added to the pixel luminance, [1, 4, 8].

In presence of a scattering medium this general scheme remains the same, only the integration spheres are now distributed over the volume, not only boundaries. Technically, the luminance integration inside the scattering medium can be performed similarly to that at the boundaries of scene objects. That is, the luminance is added when the extinction events of the forward ray is within the integration sphere [4]. The solution can be efficient in case of strong scattering, for example, for high concentration of scattering particles.

An alternative approach is gathering of the luminance when the light ray intersects the integration sphere [4]. Usually this approach provides higher calculation efficiency for most cases of scattering media, including fog and objects with low scattering properties of the media. In [9] the "point-point" and "point-beam" estimators are detailed and compared.

It is important that the luminance of the medium can be subdivided into several components: one created by the forward rays that were direct or caustic *before* entering the medium, and one created by the forward rays that were *diffuse* before entering it. In most cases, the convergence is faster if we do not calculate the latter component but instead we take luminance of the surface point where the camera ray hits diffuse surface after leaving the medium.

The integration *spheres* can be replaced by other integration volumes like e.g. cylinders which essentially affects efficiency of the method. Investigations done in [9] show that under different conditions cylinders can be advantageous over spheres or vice versa. In this paper we present another approach of "multiple integration spheres" when there can be several of them per ray segment. Comparing with cylinders, the noise for the same *number of rays* is still higher, but processing of an FMCRT ray is considerably faster than for cylinders and thus the noise after the same *time* of calculations can be reduced.

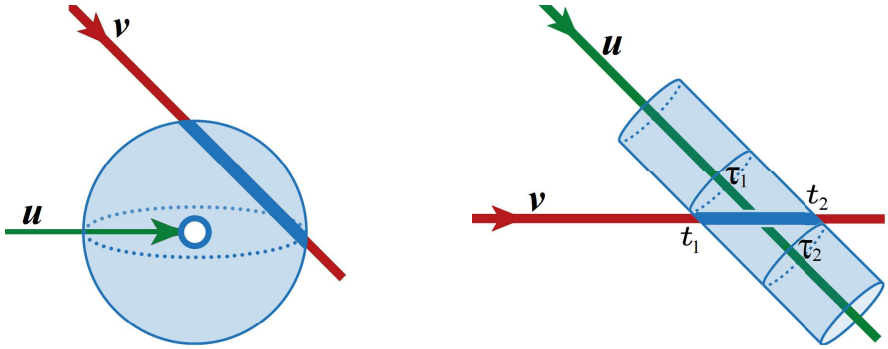

Figure 1: Integration sphere (left) and integration cylinder (right). The red arrow is a forward ray and the green arrow is a camera ray.

2 LUMINANCE OF TURBID MEDIUM IN BI-DIRECTIONAL MCRT

For view directions that enter a turbid medium, the observed radiance is the own radiance of the medium plus luminance of surfaces seen through it [10], [11]:

$$L = \underbrace{\int e^{-\sigma_{ext}t}\sigma_{sc}f(-\boldsymbol{u},\boldsymbol{v})F(\boldsymbol{v},\boldsymbol{x}_c+\boldsymbol{u}t)d^2\boldsymbol{v}dt}_{\text{medium}} + \underbrace{e^{-\sigma_{ext}T}L_s}_{\text{surface}} \tag{1}$$

where $\boldsymbol{x}_c$ is the point where the camera ray enters the medium, $\boldsymbol{u}$ is the view direction (from camera), $\boldsymbol{v}$ is illumination direction, t is the distance from $\boldsymbol{x}_c$ along the camera ray, T is the maximal t on exit from the medium, F is irradiance inside the medium and L_s is the radiance of the surface hit by the ray $\boldsymbol{x}_c + \boldsymbol{u}t$, f is the phase function and σ_{sc} is scattering of the medium and σ_{ext} is its extinction.

Here and below we use the term "luminance" as an equivalent of "radiance".

2.1 The general scheme of photon map in a turbid medium

The most straightforward variant of calculation of bi-directional MCRT in presence of turbid medium is like this. When entering the medium, the camera ray goes until the first extinction (absorption or scattering) event; then it terminates and an integration sphere is set at its end which collects the forward rays, see Figure 1.

The distance t the camera ray goes till the extinction event is random and distributed like $\sigma_{ext}e^{-\sigma_{ext}t}$. When a forward ray hits the sphere inside medium, one calculates $\sigma_{sc}f(-\boldsymbol{u},\boldsymbol{v})/(\sigma_{ext}\pi R^2)$ and adds it to the accumulated pixel luminance; here R is the sphere radius. Its average of the forward ray contribution over the FMCRT ensemble is $\int \sigma_{sc}f(-\boldsymbol{u},\boldsymbol{v})F(\boldsymbol{v},\boldsymbol{x}_c+t\boldsymbol{u})d^2\boldsymbol{v}$, and averaging over the random t just gives the integral (1), see [3, 11, 12, 10, 9].

Alternatively, we can use an integration *cylinder* instead of sphere and a deterministic integration along the camera ray instead of MC. When the camera ray enters the medium, it goes *straight* throughout the medium domain and this is the axis of the cylinder.

The contribution of a forward ray segment which intersects the cylinder

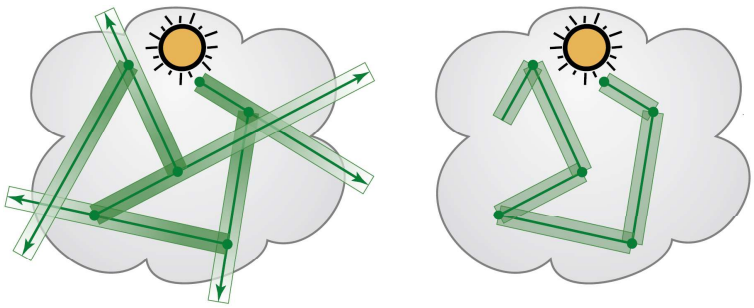

Figure 2: Long (left) and short (right) integration cylinder around the camera ray segments, from [9]. Opacity of shading shows point weight.

$$\frac{\sigma_{sc} f(\boldsymbol{u}, \boldsymbol{v})}{\sigma_{ext} \pi R^2} \frac{e^{-\sigma_{ext}\tau_1} - e^{-\sigma_{ext}\tau_2}}{|(\boldsymbol{u} \cdot \boldsymbol{v})|} \tag{2}$$

where τ_1, τ_2 are distances between the cylinder's origin and the two points of its intersection by the forward ray segment, [11], see Figure 1

As to the luminance of surfaces seen through the medium, it is taken as usual but for the rays that transmitted the medium. Their fewer number accounts for attenuation in the medium.

2.2 Shape and position of the integration volume

There can be also many other variations of this basic idea [3, 11, 12, 9]. For example, in the above method of integration cylinder the attenuation of the FMCRT ray is simulated with its random termination after passing the distance t. Instead, one can consider each FMCRT segment as a semi-infinite line and account for attenuation by a deterministic integration with weight $e^{-\sigma_{ext}t}$ along it, see Figure 2. The contribution of that "long" segment is then

$$\frac{\sigma_{sc} f(\boldsymbol{u}, \boldsymbol{v})}{\sigma_{ext} \pi R^2} \frac{e^{-\sigma_{ext}(\tau_1+t_1)} - e^{-\sigma_{ext}(\tau_2+t_2)}}{1 + |(\boldsymbol{u} \cdot \boldsymbol{v})|} \tag{3}$$

instead of (2), see [9]. Here t_1 and t_2 are the distances between the forward segment's origin and the two points where it intersects the cylinder.

A more serious alteration of the method is based on the bi-directional MCRT idea of solution of the self-consistent global illumination equation. The light field is self-consistent, i.e. the secondary illumination is created by luminance of the surfaces and media. So, instead of "sensing" the luminance in the *first* camera hit, we let the camera ray scatter there and go further, and take luminance in the subsequent hit points. To this end, integration spheres are set in those points.

Such a BMCRT step is equivalent to applying a conjugate operator to the light field, see Neumann series in [8]. But to obtain a correct result, the first integration spheres collect only direct and caustic illumination while the diffuse component is taken only in the last one, [13].

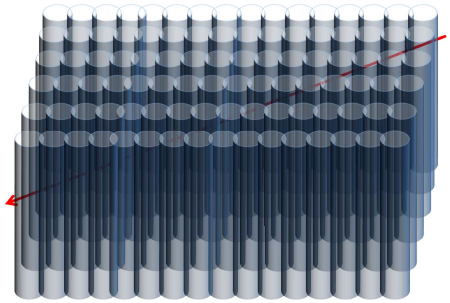

Figure 3: Forward ray intersects a lot of cylinders; shown is the simplest case of an orthogonal camera and no refraction at the medium boundary when the camera rays in the medium are parallel.

This idea can be applied to a turbid medium as well. Camera ray is traced past the first scattering event. One can either set integration spheres in all extinction points, or one can set integration cylinders along all segments of camera rays. But again, diffuse illumination is collected only by the last one, the rest collecting only direct and caustic.

On the one hand, this approach is advantageous because it increases the integration volume thus "efficiency of registration" of FMCRT rays. On the other hand, voxelization of the integration areas becomes considerably more expensive in memory. For example, in media with scattering in a narrow forward cone and high extinction coefficient, the camera ray may undergo even hundreds of scattering events before leaving the medium. Correspondingly, there are hundreds of cylinders (or spheres) per pixel.

In both these versions the whole BMCRT trajectory is used for setting integration volumes. So, if it is short, then the integration volume may be insufficient. If the trajectory is long, the integration volume can be excessive and FMCRT rays processing is expensive.

A possible solution is to separate that. The camera ray is traced long until it leaves the medium (and hits a diffuse surface). There it takes its luminance. Meanwhile the integration cylinder is along the first segment in the medium only, or, if used are integration sphere, it is set in the first extinction point [8]. The FMCRT rays scattered by diffuse surfaces are not collected by integration volumes inside the medium to avoid double counting since camera ray leaves the medium and takes luminance of diffuse surfaces. So inside the medium only the direct, caustic and scattered by the medium illumination is taken [8].

This method is very advantageous for medium with strong scattering in a narrow forward cone. In this case a thick plate of such medium is rather translucent, albeit turbid, so one can see through it. Only the last variant simulates that vision rather efficiently.

2.3 Multiple spheres as a compromise with cylinders

The main types of integration volumes are thus cylinders along the came ray segments or spheres in the ends of those segments. Cylinders provide higher efficiency of "utilization" of FMCRT rays, because they occupy large volume and it is rare that a FMCRT misses them. On the contrary, it usually intersects a lot of cylinders. For example, if we

imagine a parallel camera looking normally at a plate of a turbid material with not high extinction, the number of cylinders a typical FMCRT ray intersects is about the size of diagonal of the camera image, see Figure 3.

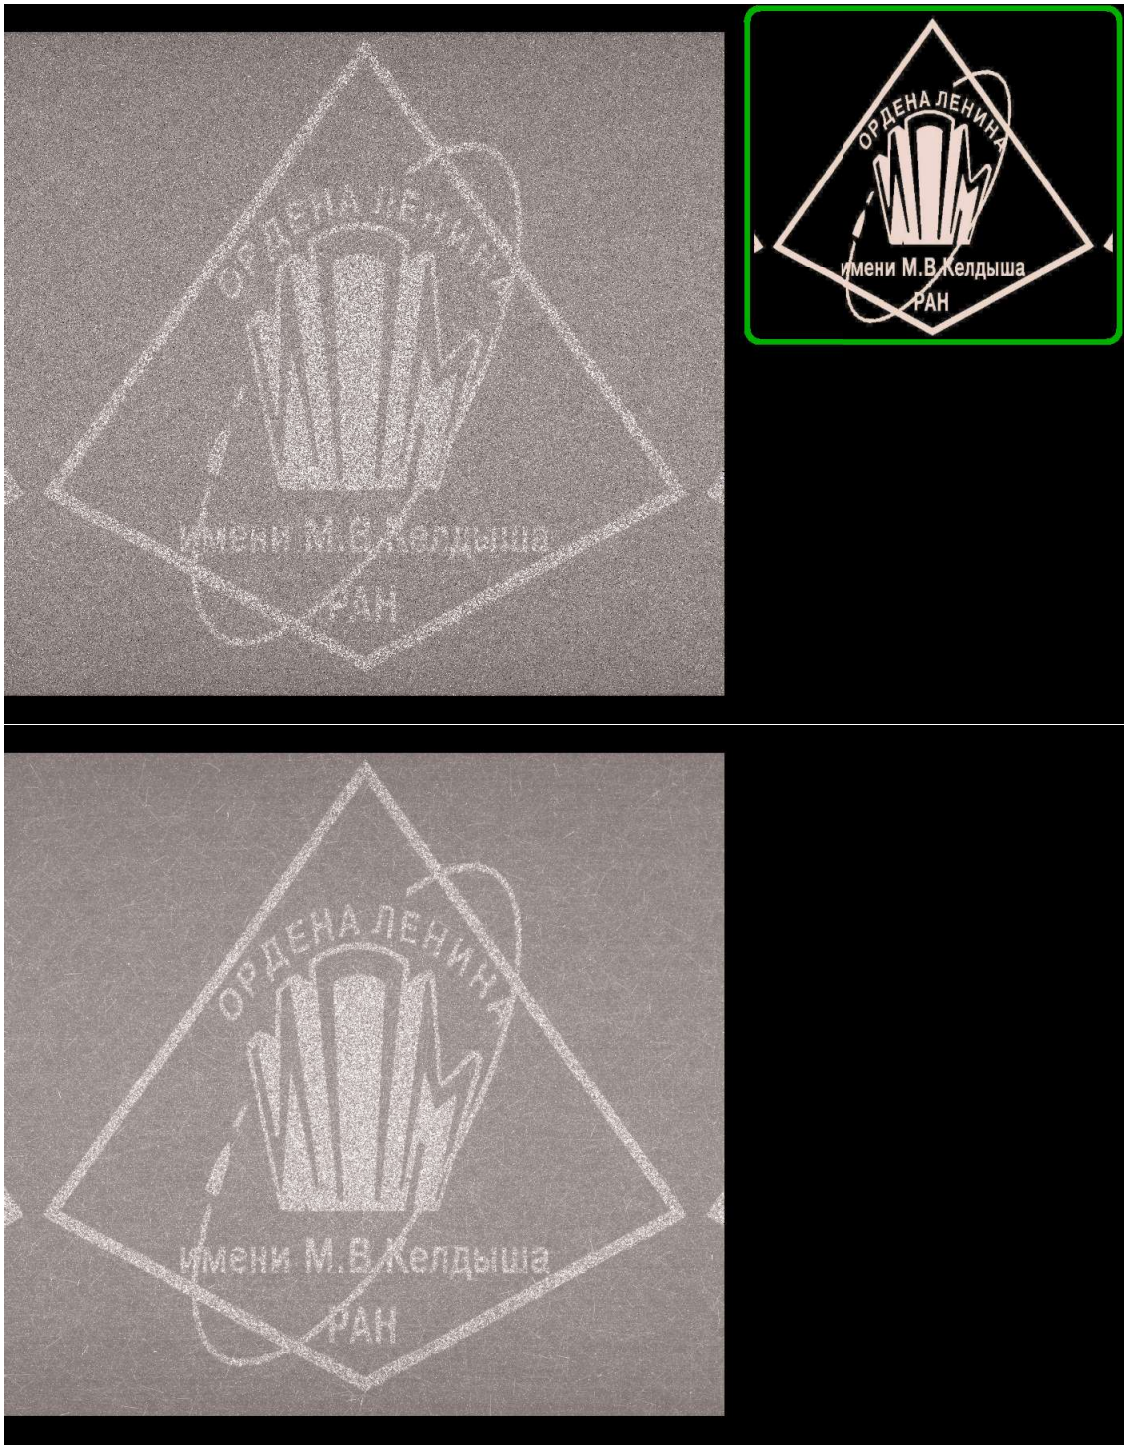

Figure 4: A cube with the textured back wall filled by scattered medium and illuminated by a parallel light. Top: single integration sphere per camera ray segment. Bottom: integration cylinder per camera ray segment. The inset in the top right corner shows the used texture (the KIAM's logo).

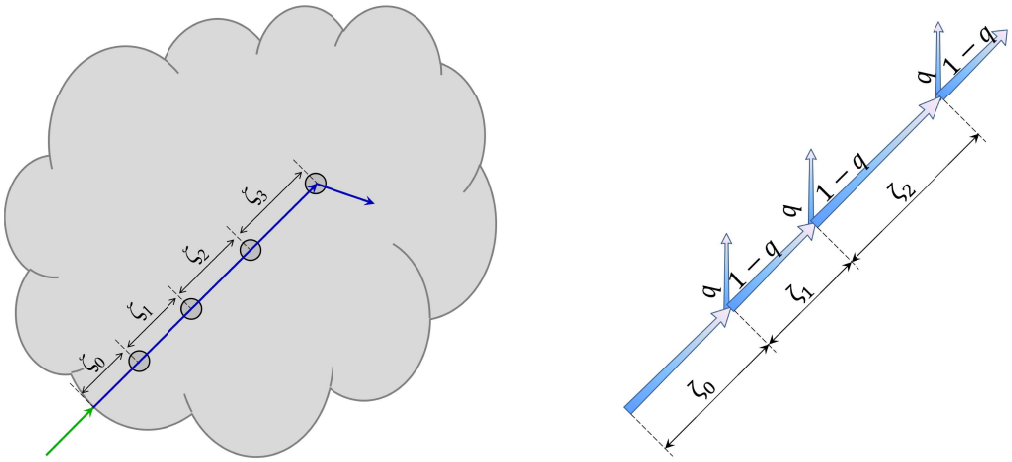

Figure 5: Several integration spheres along the camera ray segment. Left: positioning of the spheres. Right: Camera ray propagation with sub-steps: it strides by ζ_m then with probability q it undergoes extinction otherwise goes the same direction

Processing of a forward ray trajectory is then very expensive. Besides calculation of luminance for some thousand camera cylinders, one must check intersection for even many of them because in case of a curved refractive boundary those cylinders intersect and there are several (many!) of them per voxel. As a result the number of FMCRT rays processed *during the same time* is much lower which leads to the the image noise increase.

Calculations for an example scene are shown in Figure 4. The scene is a cube of 1 m size whose back wall is Lambert with albedo 40% modulated by the texture. The cube is oserved by a parallel camera whose view direction is normal to the back wall and illuminated by a parallel light inclined at 45 degrees from the view direction. The box is filled with a scattering medium with extinction coefficient 0.001 $[mm^{-1}]$ and Henyey-Grenstein phase function [14] with anisotropy parameter 0.5. The images calculated during 200 seconds for the extreme cases: single integration sphere (top) and integration cylinder (bottom) are shown in Figure 4 and one can see how the noise is lower in case of cylinders.

A possible compromise between spheres and cylinders is then to set *several* spheres per camera ray segment instead of one, see Figure 5. Increasing their number increases efficiency of "utilization" of forward rays, but in the meanwhile decelerates FMCRT part. Until there are too few intersections of FMCRT rays with the integration spheres, the deceleration effect is weak while the increase of efficiency is essential. But as their number increases so that utilization goes to 100%, the further gain is impossible while deceleration becomes essential.

This allows to find a reasonable compromise.

Additional spheres are set about the internal segment points, and one can either do that deterministically, e.g. distributing them uniformly or with exponentially decreasing density, or one can set them stochastically, choosing the distance between centers ζ at

random, see Figure 5. Here one makes a "step" of ray propagation, sets sphere, then makes another "step" and so on until the segment ends and extinction occurs, see Figure 5.

This paper describes the stochastic method. It derives the distribution of step length, criterion for extinction event (instead of propagation of the straight path) so that the accumulated luminance to converge to the correct value. Also we calculate the variance of that luminance and how it depends on the parameter that controls the (average) number of spheres.

The calculation is done for a single camera ray i.e. image pixel, because they are processed rather independently.

3 IMPLEMENTATION OF THE MULTIPLE INTEGRATION SPHERES

As suggested above, when tracing an camera ray we perform many "sub-steps" so that the ray goes straight during several of them and only then an extinction (scattering or absorption) occurs. Absorption is processed as a "Russian roulette" when the ray is killed at random. The length of sub-step ζ is random and chosen independently from the previous step so that its density is always the same $p_\zeta(\zeta)$. We assume a homogeneous isotropic medium so this density is independent from the space point.

After the step length ζ had been chosen, the ray propagates to its end and its further destiny is decided at random. With probability q there is an extinction event (scattering or absorption is then decided at random); otherwise the next step is made *retaining the ray direction*, see Figure 5.

3.1 The distribution of step length

The probability that there were n steps before extinction is therefore

$$P(n) = q(1-q)^{n-1}$$

(notice $\sum_{n=1}^{\infty} P(n) = 1$).

Density of distribution of the sum of n steps (i.e. the ray path length) is

$$p_n(s) = \int p_{n-1}(s-\zeta)p_\zeta(\zeta)d\zeta, \tag{4}$$

$$p_0(s) = p_\zeta(\zeta) \tag{5}$$

Probability of distance l before extinction is

$$P(1)p_\zeta + P(2)p_\zeta * p_\zeta + \cdots = \sum_{n=0}^{\infty} P(n+1)p_\zeta * p_\zeta * p_\zeta * \cdots * p_\zeta \tag{6}$$

(the asterisk stands for the convolution) so its Fourier image is

$$q\hat{p}_\zeta(k) + q(1-q)(\hat{p}_\zeta(k))^2 + \cdots = \frac{q\hat{p}_\zeta(k)}{1-(1-q)\hat{p}_\zeta(k)}$$

The density (6) must coincide with $\sigma_{ext}e^{-\sigma_{ext}l}$ where σ_{ext} is the extinction coefficient of medium. Therefore their Fourier images also coincide which means that

$$\hat{p}_\zeta(k) = \frac{\sigma_{ext}}{\sigma_{ext} + iqk} = \frac{1}{1 + i\frac{k}{\sigma_{ext}/q}}$$

or

$$p_\zeta(\zeta) = \alpha e^{-\alpha\zeta} \tag{7}$$

$$\alpha \equiv \frac{\sigma_{ext}}{q} \tag{8}$$

Applying (4) iteratively with this p_ζ we obtain

$$p_n(s) = \alpha e^{-\alpha s}\frac{(\alpha s)^n}{n!} \tag{9}$$

3.2 FMCRT segment contribution

Consider a single FMCRT ray segment that crosses the BMCRT ray at distance s. Its contribution is denoted $C(s)$.

The probability that there is n-th splatting sphere at s is

$$p_{n-1}(s)(1-q)^{n-1}ds = \alpha e^{-\alpha s}\frac{(\alpha(1-q)s)^{n-1}}{(n-1)!}ds$$

so the total contribution (sum over $n = 1, \ldots$) is

$$C(s)\alpha e^{-\alpha s}\sum_{n=0}^{\infty}\frac{(\alpha(1-q)s)^n}{n!}ds = C(s)\alpha e^{-\alpha s}e^{+\alpha(1-q)s}ds = \frac{C(s)}{q}\sigma_{ext}e^{-\sigma_{ext}s}ds$$

Soto get the same result as in the "standard" method (i.e. for $q = 1$) which is

$$\mathscr{C}(s)\sigma_{ext}e^{-\sigma_{ext}s}ds$$

it must be

$$C(s) = q\mathscr{C}(s) \tag{10}$$

which thus applies to both *single segment* and *whole forward ray trajectory.*

Notice the contribution of one segment of forward in the "standard" method is

$$\mathscr{C}(s) = \frac{\sigma_{sc}f(-\boldsymbol{u}, \boldsymbol{v})}{\sigma_{ext}\pi R^2}\mathcal{F} \tag{11}$$

where $\mathcal{F}$ is the total flux (sum over all light sources).

3.3 Colored case

A possible solution is that we *first* pick at random the color channel λ and then *pick ray steps according to this color channel* until extinction. Then the sum of steps until extinction has density $\sum_\lambda E_\lambda q \alpha_\lambda e^{-q\alpha_\lambda s}$ i.e. like for the standard method.

Now let us calculate contribution of a single FMCRT segment. As said above, with probability

$$\frac{E_{\lambda_0}}{\sum_\lambda E_\lambda}$$

we choose color channel λ_0. Then, like in monochrome case, the step length s has exponential density $\alpha_{\lambda_0} e^{-\alpha_{\lambda_0} s}$ where now

$$\alpha_{\lambda_0} = q\sigma_{ext,\lambda_0}$$

The probability that there is n-th splatting sphere at s is

$$p_{n-1}(s)(1-q)^{n-1}ds = \alpha_{\lambda_0} e^{-\alpha_{\lambda_0} s} \frac{(\alpha_{\lambda_0}(1-q)s)^{n-1}}{(n-1)!} ds$$

so the total contribution (sum over $n = 1, \ldots$) is

$$C_\lambda(s)\alpha_{\lambda_0} e^{-\alpha_{\lambda_0} s} \sum_{n=0}^{\infty} \frac{(\alpha_{\lambda_0}(1-q)s)^n}{n!} ds = C_\lambda(s)\alpha_{\lambda_0} e^{-\alpha_{\lambda_0} s} e^{+\alpha_{\lambda_0}(1-q)s} ds = C_\lambda(s)\alpha_{\lambda_0} e^{-q\alpha_{\lambda_0} s} ds$$

Recalling that the above is with probability $\frac{E_{\lambda_0}}{\sum_{\lambda'} E_{\lambda'}}$ we obtain the net contribution

$$q \frac{\sum_{\lambda_0} E_{\lambda_0} \sigma_{ext,\lambda_0} e^{-\sigma_{ext,\lambda_0} s}}{\sum_\lambda E_\lambda} C_\lambda(s) ds$$

In the "standard" method (i.e. for $q = 1$) it is

$$\frac{\sum_{\lambda_0} E_{\lambda_0} \sigma_{ext,\lambda_0} e^{-\sigma_{ext,\lambda_0} s}}{\sum_\lambda E_\lambda} \mathscr{C}_\lambda(s) ds$$

so for both single segment and multi-segment FMCRT path, the contribution is

$$C_\lambda(s) = q\mathscr{C}_\lambda(s)$$

4 CONCLUSION

We produced a method of "gathering photons" in scattering medium which uses several integration spheres stochastically distributed over a camera ray segment. This forms a disjoint integration volume which is a compromise between usual integration sphere (one per segment) and integration cylinder. Changing the control parameter q that determines the mean number of spheres, we can gradually change from single integration spheres to a cylinder (a union of a large number of spheres).

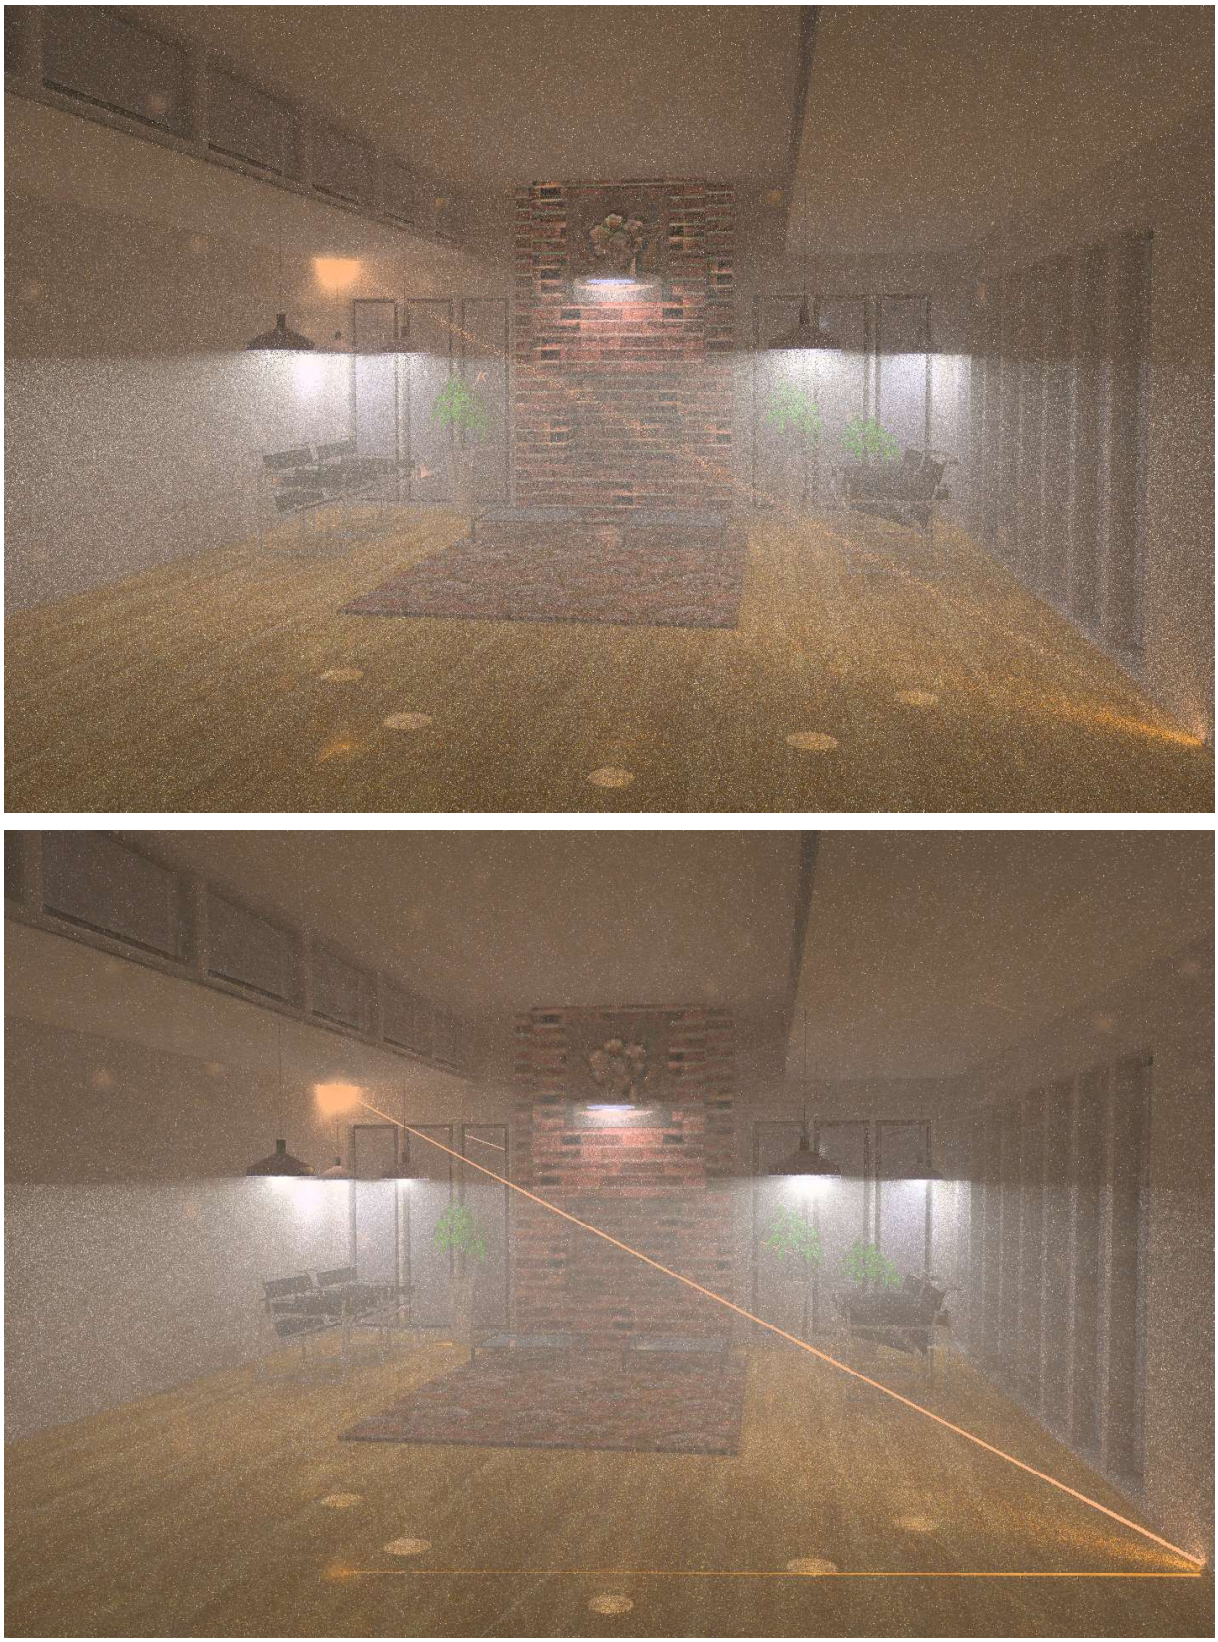

Figure 6: A room filled by fog illuminated by point light sources and laser beam from the right bottom corner to the top left corner. Top: single integration sphere per camera ray segment. Bottom: integration cylinder per camera ray segment.

When choosing the optimal control parameter of the method (or, in other cases, deciding between various other integration volumes) one must realize that those that provide lowest noise of the medium luminance for the fixed number of rays are not always really the best. This is because, first, some advance methods with integration volumes occupying most of the medium domain seriously decelerate ray tracing. As a result, the number of rays processed during the same *time* of calculations can drop so seriously that the noise after the fixed simulation time is *higher* (as compared to competing methods). Second, the own luminance of the medium is not the only component; there is also the luminance of objects "seen through" that medium. If both are calculated from the same ray sets it may happen that those optimal for the own luminance are bad for the luminance seen through me medium, and vice versa.

Calculations for an example scene of a room filled with light "fog" are shown in Figure 6. The room is illuminated by several point lights sources under the conical shades and besides there is a thin laser beam from the right bottom corner to the top left corner. This floor is slighly specular so we can see reflections of the lamps and the laser beam in it. The "fog" has extinction coefficient 0.0001 $[mm^{-1}]$ and Henyey-Grenstein phase function [14] with anisotropy parameter 0.5. The images calculated during 4000 seconds for the extreme cases: single integration sphere (top) and integration cylinder (bottom) are shown in Figure 6.

One can see how the noise is lower for integration cylinders and also this method correctly reproduces the *thin* laser beam while for single integration sphere it is made unnaturally thick and consists of sparse dots, and at last we do not see its reflection in the floor. The noise see in the bottom image of Figure 6 is mainly from the diffuse *surfaces* while the radiance of the fog is nearly smooth.

Analysis of noise with due account of these factors is our future work plan.

REFERENCES

[1] E. P. Lafortune and Y. D. Willems, "Bi-directional path tracing", in *Proceedings of Third International Conference on Computational Graphics and Visualization Techniques (Compugraphics '93), Alvor, Portugal, December 1993*, 145–153 (1993).

[2] C. Hery, R. Villemin, and F. Hecht, "Towards bidirectional path tracing at pixar", in *Physically Based Shading in Theory and Practice, ACM SIGGRAPH 2016 Courses* (2016). Available: http://graphics.pixar.com/library/BiDir/

[3] W. Jarosz, M. Zwicker, and H.W. Jensen, "The Beam Radiance Estimate for Volumetric Photon Mapping", *Computer Graphics Forum*, **27** (2), 557–566 (2008).

[4] D.D. Zhdanov, A.A. Garbul, I.S. Potemin, A.G. Voloboy, V.A. Galaktionov, S.V. Ershov, and V.G. Sokolov, "Photorealistic volume scattering model in the bidirectional stochastic ray tracing problem", *Programming and Computer Software*, **41** (5), 295–301 (2015). https://doi.org/10.1134/S0361768815050102

[5] T. Hachisuka and H.W. Jensen, "Stochastic progressive photon mapping", *ACM Trans. Graph.*, **28** (5), 141:1–141:8 (2009). http://doi.acm.org/10.1145/1618452.1618487

[6] V. Havran, R. Herzog, and H.-P. Seidel, "Fast Final Gathering via Reverse Photon Mapping", *Computer Graphics Forum*, **24** (3), 323–333 (2005).

[7] H.W. Jensen, “Global illumination using photon maps”, in *Proceedings of the Eurographics Workshop on Rendering Techniques ’96*, 21–30 (1996). http://dl.acm.org/citation.cfm?id=275458.275461

[8] S.V. Ershov, D.D. Zhdanov, A.G. Voloboy, and M.I. Sorokin, *Treating diffuse elements as quasi-specular to reduce noise in bi-directional ray tracing*, Preprint IPM No. 122 (Moscow: KIAM), (2018). doi:110.20948/prepr-2018-122-e

[9] J. Křivánek, I. Georgiev, T. Hachisuka, P. Vévoda, M. Šik, D. Nowrouzezahrai, and W. Jarosz, “Unifying points, beams, and paths in volumetric light transport simulation”, *ACM Trans. Graph.*, **33** (4), 1–13, (2014).

[10] E.P. Lafortune and Y.D. Willems, “Rendering participating media with bidirectional path tracing”, in *Rendering Techniques ’96*, 91–100 (1996).

[11] W. Jarosz, D. Nowrouzezahrai, I. Sadeghi, and H.W. Jensen, “A comprehensive theory of volumetric radiance estimation using photon points and beams”, *ACM Trans. Graph.*, **30** (1), 5:1–5:19 (2011). http://doi.acm.org/10.1145/1899404.1899409

[12] W. Jarosz, D. Nowrouzezahrai, R. Thomas, P.-P. Sloan, and M. Zwicker, “Progressive photon beams”, *ACM Trans. Graph.*, **30** (6), 181:1–181:12 (2011). http://doi.acm.org/10.1145/2070781.2024215

[13] H.W. Jensen and P. Christensen, “High quality rendering using ray tracing and photon mapping”, in *ACM SIGGRAPH 2007 Courses*, ser. SIGGRAPH ’07. (2007). http://doi.acm.org/10.1145/1281500.1281593

[14] G. Pomraning, “On the Henyey-Greenstein approximation to scattering phase functions”, *Journal of Quantitative Spectroscopy and Radiative Transfer*, **39** (2), 109 – 113 (1988). http://www.sciencedirect.com/science/article/pii/0022407388900787

Received January 18, 2019