Light & Engineering
Vol. 28, No. 5, pp. 39–47, 2020
https://doi.org/10.33383/2019-103
Svetotekhnika #3, 2020, pp. 64–70

THE METHOD OF QUASI-SPECULAR ELEMENTS TO REDUCE STOCHASTIC NOISE DURING ILLUMINANCE SIMULATION

Sergei V. Ershov[1,1], Dmitry D. Zhdanov[2,2], Alexei G. Voloboy[1,3], and Nikolay B. Deryabin[1,4]

[1]Department of Computer graphics and computational optics of the Keldysh Institute of Applied Mathematics RAS, Moscow, Russia
[2]ITMO University, St. Petersburg, Russia
E-mail: measure@spp.keldysh.ru[1], ddzhdanov@mail.ru[2], voloboy@gin.keldysh.ru[3], dek@gin.keldysh.ru[4]

ABSTRACT

When simulating the propagation of light, luminance/radiance brought by a ray is calculated from the optical properties of the scene objects it interacts with. According to their optical properties, objects can be roughly divided into diffuse and specular. In Monte Carlo ray tracing luminance/radiance is calculated only for diffuse surfaces. When a ray hits a specular a surface, it is reflected (or refracted) until it reaches a diffuse surface, and only then the luminance/radiance is calculated. In the proposed approach, diffuse elements are further divided into genuine diffuse and quasi-specular elements. The most natural criterion for the latter is that it scatters light in a narrow cone about the specular direction. An element of the scene can also be a superposition of both types when its scattering function is a sum of the genuine diffuse and quasi-specular parts. This article shows how different components of illuminance/irradiance interact with quasi-specular objects and describe how this works in the bi-directional stochastic ray tracing. The proposed approach significantly reduces stochastic noise for multiple scenes. This method is also applicable for simulation of volume scattering, treating the phase function of the medium as quasi-specular. In this case, the choice of quasi-specular objects is not based on the nature of the bidirectional scattering distribution function (BSDF): the medium is treated as completely quasi-specular while the surfaces, even if their BSDFs are narrower, remain genuine diffuse. The article shows the advantage of this approach.

Keywords: calculation of illuminance, realistic rendering, bidirectional ray tracing, stochastic ray tracing, noise reduction, BSDF

1. INTRODUCTION

The bidirectional ray tracing using photon maps is an efficient method for calculating the image of a virtual scene [1]. Tracing rays from light sources creates a photon map that allows one to calculate the illuminance of the scene surfaces. Then the ray is traced from the camera, and when it hits a diffuse surface, it takes illuminance from the photon map, “convolves” it with the surface bidirectional scattering distribution function (BSDF) at this point, and adds the result to the accumulated luminance/radiance of the pixel. This idea is implemented in several modifications of the described method [2–5].

An important parameter of this approach is the number of operators (events on the ray path) in the products of integral scattering operators calculated by the backward Monte Carlo ray tracing. Usually only diffuse events are considered, and the maximum allowed number is denoted by backward diffuse depth (*BDD*[1]). The efficiency of the method, that is, its convergence rate (or noise level, which is

[1] A specific parameter of bidirectional ray tracing: the ray from camera terminates after BDD diffuse scattering events.

the same) strongly depends on this *BDD*, and its optimal value differs for each scene.

The best approach would be to use different *BDDs* for different parts of the scene [5], and even mix calculations with different *BDDs* [5], as in "multiple importance sampling". It is often possible to find the optimal *BDD* (automatically or manually), for which calculations become quite efficient. However, this is not always possible, and in some cases, changing the *BDD* does not help, i.e. whatever value is taken, the image is very noisy.

We provide a simple method that helps in such cases. Its idea is that those objects of the scene (surface or volume scattering) that have a diffuse part of BSDF are divided into two groups: genuine diffuse and quasi-specular. Usually, the last ones will be those for which the BSDF is a narrow near-specular cone. Although in principle the criterion can be arbitrary, up to the point that the Lambert surface will be treated as quasi-specular. As a rule, the separation treats the BSDF (for the surface or the phase function for a volume) as a sum of the quasi-specular and genuine diffuse components, which are both non-zero.

Quasi-specular scattering does not increase the diffuse event counter (when it exceeded the *BDD*, the ray is killed). Besides, the quasi-specular part of the BSDF does not convolve with diffuse component of illuminance (that is, those rays from the light source that have been subjected to genuine diffuse scattering). This change in backward ray tracing can be applied to both surface and volume scattering, to reduce noise and the amount of memory used.

An alternative approach is more attractive, in which there is no fixed *BDD*, no a distinction between direct, caustic and diffuse lighting rays. At each intersection of a diffuse surface by the camera ray, all the lighting rays are collected, and the contribution from various points is summed with the weight depending on the full trajectory (joining its parts from the light source and from the camera) and calculated using the multiple importance sampling (MIS) equation [6]. Unfortunately, despite the convincing images in [6], the proposed method does not use photon maps at all, but instead traces one ray from the camera and one ray from the light source, then calculates the contribution from the joining of these paths (it vanishes if they do not meet with the desired accuracy), then forgets this pair and starting a new one. When using photon maps, the same forward ray path is tried to join with each ray from the camera. At first glance, this is all the same since contributions for different forward paths are independent even for the same backward path. In fact, this is not the case [7], and the noise (that is, the variance of the accumulated pixel luminance) is calculated using equations other than those used for the simple Markov chain Monte-Carlo (MCMC). These equations include not only the BSDF along the ray path and the light source goniogram, as in [6], but the geometric factors (distance between surfaces, etc.) as well. Therefore, the method proposed in [6] cannot be used directly without improvement. Meanwhile, it requires that integrating spheres have been installed at all hit points for the camera ray, which increases the memory requirements. And in addition, the photon (the ray from the light source) is checked for intersections with all of them, which slows down the processing process. Therefore, the performance of our proposed method based on a limited *BDD* is higher.

There are variants of the method [6], gradually abandoning photon maps, for example [8], wherein addition to vertex merging, which is much similar to photon maps, the vertex connection is used when the ends of the camera and light paths are not close and connected by an additional segment of finite length. Vertexes for joining can be selected probabilistically [9]. A detailed review can be found in the dissertation [10]. All these methods use MIS, when the merging of the camera and light paths is done at different vertexes and summing the resulting contributions with weights. In addition, in [8] MIS is applied to two possible approaches – vertex connection and vertex merging.

Unfortunately, as in [6], in these studies, the equations for weight calculations are applied to photon maps or their equivalent when the same set of paths from the source is reused for all camera rays.

2. THE RENDERING EQUATION AND CALCULATION OF ITS NEUMANN SERIES BY BIDIRECTIONAL RAY TRACING

The light field in the scene is described by a self-consistent equation that can be written in various forms, such as the equation of global illumination [11], etc. The idea is that there is a light field in the scene, it lights the surfaces that scatter it. This is the transformation of irradiance E (light incident on the surface at a point $\boldsymbol{x}$ in direction $\boldsymbol{v}'$) into luminance/radiance L (light exiting from a surface point $\boldsymbol{x}$ in direction $\boldsymbol{v}$) by the BSDF of the surface f:

$$L(v;x) = \int f(x;v,v')E(v',x)d^2v'.$$

Then this light emitted from the surface propagates further across the scene and lights its surfaces. This transformation of luminance/radiance (light emitted from the surface) into illuminance (light incident on another point in the scene) is described by the transfer operator. Since the explicit form of the scattering and transport operators is not important for us, we will use compact notation:

$$L = \hat{F} \cdot E, \tag{1}$$

$$E = \hat{T} \cdot L. \tag{2}$$

Note that here $\hat{F}$ includes only diffuse scattering; therefore, the caustic is non-scattered light, that is, part of direct lighting.

Note that L is the luminance/radiance of scene surfaces, while the camera image will be $\hat{S}L$, where $\hat{S}$ describes a purely specular transformation between the scene and the camera (usually it is identity operator). Here and further, we will calculate only L. Then, if necessary, we can apply its conversion to the luminance of the image.

Full illuminance consists of direct and diffuse components:

$$E = E^{(i)} + E^{(0)}. \tag{3}$$

By combining these three equations, we arrive at the self-consistent global illumination equation, used below in the form of the rendering equation introduced by Kajiya [11]:

$$E = \hat{T} \cdot \hat{F} \cdot E + E^{(0)}. \tag{4}$$

In computational optics a combination of forwarding and backward ray tracing is widely used, when the forward part calculates illuminance of diffuse surfaces E and stores it, for example, as a photon map [1], and then the backward tracing converts it into an image visible by the camera. Note that the illuminance obtained by stochastic tracing is noisy, and this noise goes further to the image, its final amplitude strongly depends on how the tracing from the camera works.

The easiest way is to trace the rays from the camera to the first diffuse surface, where the surface luminance/radiance created by all components of the illuminance is calculated ($= \hat{F}E$), which is then added to the accumulated luminance/radiance of the pixel. This provides an estimate of the luminance/radiance of the surface (1), although the result is not ideal, since the illuminance E calculated by the forward ray tracing is usually subjected to a (spatial) filtering to reduce noise. Because of this, small-scale lighting details, such as glare, can be lost.

Instead, one can apply the Nth iteration of (4) to illuminance, which leads to the luminance/radiance equation:

$$L = (\hat{F} \cdot \hat{T})^N \hat{F} \cdot E + \sum_{k=0}^{N-1} (\hat{F} \cdot \hat{T})^k \hat{F} \cdot E^{(0)}. \tag{5}$$

For the exact illuminance E, this naturally gives the same result as the simple $\hat{F}E$ above, however for the actually used noisy E, this second form is often better due to the lower noise level due to the convolution with the power of operator $\hat{F}\hat{T}$. The value N is nothing but the "backward diffuse depth" or *BDD* described in the introduction.

The term $\hat{F}E^{(0)}$ is the luminance/radiance of the surface under direct (including caustic) lighting. The diffuse component of the lighting is considered only at the last hit point of the camera ray: this is the term $\left(\hat{F}\hat{T}\right)^N \hat{F}E$.

Integral operators can be calculated by the Monte Carlo method: we launch rays from the camera through a given pixel of the image, they hit a surface, are scattered (left operator $\hat{F}$), spread across the scene (operator $\hat{T}$; note that in the backward ray tracing the order of events from camera corresponds to the order of operators from left to right, i.e. the leftmost operator corresponds to the transformation in the segment closest to camera), and so on until the Nth diffuse surface is reached, where it terminates. At the kth diffuse hit point (i.e. immediately before the kth diffuse scattering), we calculate the luminance/radiance of the surface under the direct (including caustic!) lighting $\hat{F}E^{(0)}$ if $k < N$, and under the full lighting $\hat{F}E$ for $k = N$. Then the result is scaled to consider the attenuation of light due to specular transformations in $\hat{T}$ and is eventually added to the accumulated luminance/radiance of the pixel. The average value over the ensemble of camera rays converges to L.

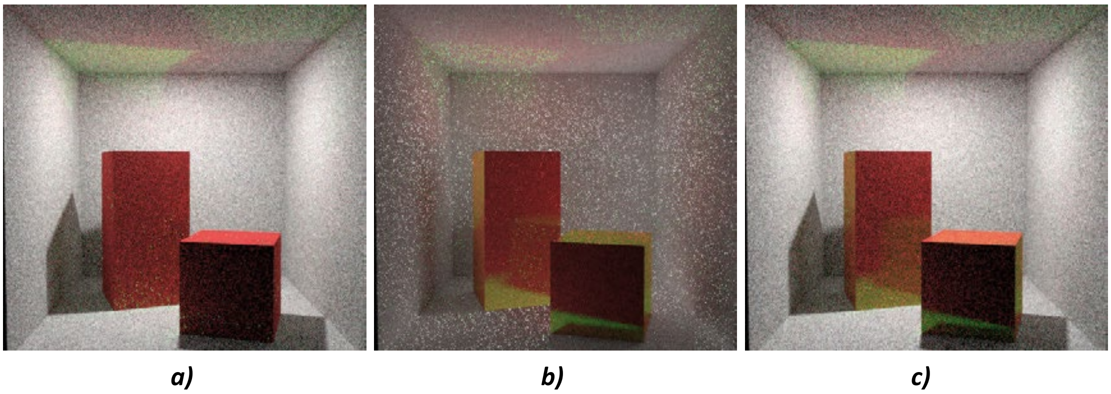

Fig. 1. Camera images for the modified Cornell Box scene calculated during the same time (200 s) for the standard method with *BDD* = 0 (a), the standard method with *BDD* = 1 (b) and for the quasi-specular method with *BDD* = 0 (c)

3. THE IDEA OF THE QUASI-SPECULAR METHOD

Consider a modification of the well-known Cornell Box benchmark scene. It is a box open from the camera side with two parallelepipeds standing on its floor. The light source is a transparent square located very close to the ceiling of the box, which emits upwards. The walls of the box have a Lambertian reflection with an albedo of 50 %, and the parallelepipeds have BSDF, which is a sum of the Lambertian and narrow (almost specular) components. To distinguish their contributions to the image, the Lambert component is pure red (the normalized colour is (1,0,0)), and the sharp component is almost pure green (the normalized colour is (0.1,1,0.1)). The albedo for the maximum colour component is 50 % and 25 %, respectively.

There is practically no direct lighting in the scene (except for a small square of the ceiling above the source).

Consider the case of *BDD* = 0. A ray from the camera that hits any surface of the scene terminates there (since there are no specular BSDFs in the scene) and takes the luminance/radiance of the hit point as a convolution of the BSDF with full lighting. There is a sharp BSDF on the surface of the parallelepipeds, while the lighting is mainly secondary from the walls of the box. The convolution of a sharp BSDF with wide noise lighting from photon maps contains strong noise (green reflection at the bottom of the parallelepipeds in Fig. 1).

In case of *BDD* = 1, only the direct and caustic lighting, which is not present in this scene, except for a small square on the ceiling above the source, is taken at the first hit point of a camera ray. Then the camera ray is reflected and either leaves the scene or hits one of its surfaces, where full lighting collects. In this case, a very small part of the rays hits the square of the ceiling located directly above the source, where the lighting is orders of magnitude higher. Accordingly, the luminance/radiance brought by these rays to the pixel is orders of magnitude higher than the average value. Other rays hit the walls of the box or parallelepipeds, where the lighting is low, and accordingly, their contribution to the pixel luminance/radiance is also small. As a result, the luminance/radiance of the pixel is created mainly by a small fraction of rays, which inevitably leads to a large amount of noise, see Fig. 1. Note that the noise on the parallelepipeds is mostly not green, because the sharp BSDF part sends reflected rays to about the same place in the scene, so either almost all of them hit the bright square on the ceiling or none.

Higher *BDD* values also do not improve the situation.

One could try to use *BDD* = 0 for camera rays reflected by the Lambert part of the BSDF, and *BDD* = 1 for the rays that were first scattered by the sharp BSDF part. This, however, does not work [5], as was proved in [12]. That is when rays scattered in the near- specular direction and off-specular ones collect lighting differently, the accumulated luminance/radiance is incorrect. Thus, a more complex criterion is required to determine at which point to use which lighting rays, and which rays to ignore.

The key idea of the approach remains the same: scattering from the sharp BSDF part is treated not as genuine diffuse, but rather like specular. This sharp BSDF (and the entire method of processing the camera rays scattered by it) is therefore called quasi-specular. The result of the calculation for the same model scene using this method is shown in Fig. 1 (c), see section 7.1. for a more detailed discussion.

4. OPERATOR SERIES IN PRESENCE OF QUASI-SPECULAR BSDFS

Now let us go to the formal derivation of what to do with the camera ray when the diffuse BSDF is divided into genuine diffuse and quasi-specular parts:

$$\hat{F} = \hat{F}_d + \hat{F}_{qs}. \quad (6)$$

The separation is arbitrary (although some separations are advantageous in terms of noise level and some are not), meaning that for any choice, the image luminance/radiance converges to the exact value.

Considering (6) the luminance at the point of the surface $\boldsymbol{x}$ is

$$L(v,x) = (\hat{F}_d \cdot E)(v,x) + (\hat{F}_{qs})(v,x). \quad (7)$$

Lighting is now also divided into three components: direct (that was not scattered at all or was purely specular BSDF), quasi-caustics (scattered at least once by a quasi-specular BSDF, but never by genuine diffuse BSDF), and diffuse (scattered at least once by genuine diffuse BSDF), i.e.

$$E = E^{(0)} + E^{(qc)} + E^{(i)},$$

so (7) takes the form

$$L = \hat{F} \cdot E = \hat{F}_{qs} \cdot (E^{(0)} + E^{(qc)}) + \\ + \hat{F}_{qs} \cdot E^{(i)} + \hat{F}_d \cdot E. \quad (8)$$

Then, substituting our light separation into the global illumination equation (4), we obtain

$$E^{(qc)} + E^{(i)} = \hat{T} \cdot \hat{F}_{qs} \cdot (E^{(0)} + E^{(qc)}) + \\ + \hat{T} \cdot \hat{F}_{qs} \cdot E^{(i)} + \hat{T} \cdot \hat{F}_d \cdot E.$$

The term $\hat{T} \cdot \hat{F}_{qs} \cdot E^{(i)} + \hat{T} \cdot \hat{F}_d \cdot E$ in its right-hand side describes light that has undergone at least one genuine diffuse scattering, while the term $\hat{T} \cdot \hat{F}_{qs} \cdot (E^{(0)} + E^{(qc)})$ describes the light that has undergone at least one purely specular scattering, but no one diffuse scattering. Considering our splitting of lighting into three components, this means that

$$E^{(i)} = \hat{T} \cdot (\hat{F}_d \cdot E + \hat{F}_{qs} \cdot E^{(i)}), \quad (9)$$

$$E^{(qc)} = \hat{T} \cdot (\hat{F}_{qs} \cdot E^{(qc)} + \hat{F}_{qs} \cdot E^{(0)}), \quad (10)$$

which means that

$$E^{(i)} \Big| = (1 - \hat{T} \cdot \hat{F}_{qs})^{-1} \cdot \hat{T} \cdot \hat{F}_d \cdot E, \quad (11)$$

$$E^{(qc)} \Big| = (1 - \hat{T} \cdot \hat{F}_{qs})^{-1} \cdot \hat{T} \cdot \hat{F}_{qs} \cdot E^{(0)}. \quad (12)$$

We assume that if the camera ray is subjected to quasi-specular scattering, this does not increase the counter of diffuse events, so the ray does not terminate. Now we will derive which illuminance components should be taken at which hit points of the camera ray so that the mathematical expectation of image luminance/radiance coincides with the exact value.

By combining (8) with (11) and (12), after trivial, though tedious transformations, it can be obtained that the surface luminance/radiance calculated for $BDD = N$ is equal to $L \Big| = \hat{F}_{qs} \cdot (E^{(0)} + E^{(qc)}) +$

$$\Big| + \sum_{k=0}^{N-1} ((1-\hat{Q})^{-1} \cdot \hat{F}_d \cdot \hat{T})^k \cdot (1-\hat{Q})^{-1} \times \\ \times \hat{F}_d \cdot (E^{(0)} + E^{(qc)}) + (1-\hat{Q})^{-1} \times \\ \times \hat{F}_d \cdot \hat{T})^N \cdot (1-\hat{Q})^{-1} \hat{F}_d \cdot E), \quad (13)$$

where

$$\hat{Q} \equiv \hat{F}_{qs} \cdot \hat{T}. \quad (14)$$

There is also an alternative form that gives the same result for the exact luminance, while for a real noisy illuminance E may give different (in noise level) result:

$$L = (1-\hat{Q})^{-1} \cdot \hat{F}_{qs} \cdot E^{(0)}$$

$$\Big| + \sum_{k=0}^{N-1} ((1-\hat{Q})^{-1} \cdot \hat{F}_d \cdot \hat{T})^k \cdot (1-\hat{Q})^{-1} \cdot \hat{F}_d \times \\ \times (E^{(0)} + E^{(qc)}) + (1-\hat{Q})^{-1} \times \\ \times \hat{F}_d \cdot \hat{T})^N \cdot (1-\hat{Q})^{-1} \hat{F}_d \cdot E). \quad (15)$$

The detailed conclusion is given in [12].

5. INTEGRATION BY CAMERA PATHS

By decomposing $(1-\hat{Q})^{-1}$ from (13) into the Neumann series, we see that, for example, for $BDD = 2$

$$L = \hat{F} \cdot (E^{(0)} + E^{(qc)}) + \sum_{k=1}^{\infty} \hat{Q}^k \cdot \hat{F}_d \times \\ \times (E^{(0)} + E^{(qc)}) + \sum_{k,m=0}^{\infty} \hat{Q}^k \cdot \hat{F}_d \cdot \hat{T} \cdot \hat{Q}^m \times \\ \times \hat{F}_d \cdot (E^{(0)} + E^{(qc)}) + \sum_{k,m,n=0}^{\infty} \hat{Q}^k \cdot \hat{F}_d \cdot \hat{T} \times \\ \times \hat{Q}^m \cdot \hat{F}_d \cdot \hat{T} \cdot \hat{Q}^n \cdot \hat{F}_d \cdot (E^{(0)} + E^{(qc)}). \quad (16)$$

In this expression, a term of the form $\hat{Q}^k \cdot \hat{F}_d \cdot \hat{T} \cdot \hat{Q}^m \cdot \hat{F}_d \cdot (E^{(0)} + E^{(qc)})$ means that:

1. The final (the last "before" camera) light transformation is k ≥ 0 quasi-specular transformations with the corresponding transfer $\hat{Q}^k = (\hat{F}_{qs} \cdot \hat{T})^k$;

2. Prior to that (i.e. further from the camera), the light undergoes a genuine diffuse transformation with the corresponding transfer $\hat{F}_d \cdot \hat{T}$;

3. Prior to that (i.e. even further from the camera), the light had undergone $m \geq 0$ quasi-specular transformations with the corresponding transfer $\hat{Q}^m = (\hat{F}_{qs} \cdot \hat{T})^m$;

4. And all this affects the convolution of the genuine diffuse BSDF component with direct and quasi-caustic lighting, i.e. on $\hat{F}_d \cdot (E^{(0)} + E^{(qc)})$.

The action of the integral operators $\hat{F}_d$ and $\hat{F}_{qs}$ can be calculated by the Monte Carlo method tracing rays from the camera. Here is the first ray transformation (along the path from the camera) corresponds to the leftmost operator in the product, and the last transformation corresponds to the rightmost operator. Thus, our term $\hat{Q}^k \cdot \hat{F}_d \cdot \hat{T} \cdot \hat{Q}^m \cdot \hat{F}_d \times \times (E^{(0)} + E^{(qc)})$ is estimated from camera rays that first underwent $k \geq 0$ quasi-specular events, then one genuine diffuse, then $m \geq 0$ quasi-specular, and then took the luminance/brightness of the genuine diffuse BSDF part under the direct and quasi-caustic lighting (i.e. diffuse lighting is ignored).

A detailed derivation is given in [12].

The other terms of equation (16) can be similarly calculated by the Monte Carlo ray tracing from the camera. This leads to the algorithm of processing camera rays in which the luminance/radiance of the hit point is calculated as follows:

- Before (and including!) the first non-purely specular event – $\hat{F} \cdot (E^{(0)} + E^{(qc)})$;
- After the first quasi-specular event and until the second genuine diffuse – $\hat{F}_d \cdot (E^{(0)} + E^{(qc)})$;
- After the second genuine diffuse event – $\hat{F}_d \cdot E$;
- And at the third genuine diffuse event the ray terminates.

This relates to *BDD* = 2. Similarly, the case of another *BDD* is considered, as well as the alternative form (13). Trace the rays from camera until they are subjected to *BDD* + 1 genuine diffuse event (or are absorbed earlier); after that, the ray terminates. When the ray hits a surface that has a diffuse or quasi-specular BSDF, it takes a convolution of a part of BSDF with a part of illuminance:

- Diffuse lighting – always only with a diffuse remainder of the BSDF;
- Quasi-caustic lighting: the main variant is before the first quasi-specular event– with a complete BSDF, after it – with the diffuse remainder of the BSDF or the alternative variant is always with the diffuse remainder of the BSDF;
- Direct and caustic lighting: the main variant is before the first quasi-specular event – with a complete BSDF, after it – with the diffuse remainder of the BSDF or the alternative variant is before the first diffuse scattering – with a complete BSDF, after it – with the diffuse remainder of the BSDF.

6. VOLUMETRIC SCATTERING

6.1. Standard Method

It works the same way as for surfaces. For simplicity, suppose that *BDD* = 1 and the camera is inside the scattering medium. Then the camera ray propagates in the medium. When it undergoes the first volume scattering, we collect direct and caustic components of illuminance and convolve it with the phase function. At the point of the second volume scattering, full illuminance is taken. After the second diffuse scattering, the camera ray terminates.

Consider now a model scene in which the camera looks through a layer of scattering medium at an illuminated diffuse surface. Let the medium also be non-absorbing and have a large scattering coefficient, so that the camera ray undergoes many scattering events in it until it reaches this surface. And finally, let the phase function be sharp so that each scattering changes the ray direction only slightly.

In the standard method with a small *BDD,* the camera ray terminates only near its entrance to the medium, so that it does not reach the diffuse surface behind the medium at all. Besides, since the phase function is narrow, its convolution with the illuminance coming from the surface of the scene leads to strong noise, see section 3.

Therefore, to reduce noise, it is necessary to use a large *BDD* so that the camera ray penetrates the medium and reaches the surface behind it, rather than collecting diffuse lighting at the points of volume scattering. However, this requires storing all the volume scattering events along the long ray path (since direct and caustic lighting are collected there), which usually requires too much memory.

6.2. Quasi-Specular Medium

Using a quasi-specular approach can significantly improve the situation for scenes of that kind.

Suppose that the surface of the scene is not quasi-specular (that is, it does not have a quasi-specular component that is processed as described at the end of section 5). Also, let the entire phase function be treated as quasi-specular, i.e. its genuine diffuse part vanishes: $\hat{F}_d = 0$. Then, while the camera ray propagates inside the scattering medium, it only undergoes quasi-specular events that do not increase the counter of diffuse events, and as a result, it propagates in the medium until it leaves it and therefore reaches the surface behind the medium.

Diffuse lighting now consists of the light reflected from the diffuse surface of the scene, quasi-caustics is the light that has undergone at least one diffuse event (and an arbitrary number of specular scattering – but not diffuse surface scattering).

Since $\hat{F}_d = 0$, the "main variant" (see the end of section 5) assumes that at the points of volume scattering:

• Diffuse lighting (from the scene surface) is ignored;

• Direct, caustic, and quasi-caustic lighting is used only at the point of the first volumetric scattering.

As a result, for any *BDD* the camera ray leaves the medium and reaches the surface of the scene. In this case, only one (first) volume scattering is remembered, and the luminance/radiance of the medium is not taken at the subsequent points. There is also no strong noise from the convolution of the sharp phase function in the diffuse component of illuminance (from the surface). However, there is its convolution with the direct, caustic and scattered inside the medium light. But usually, they do not have a wide angular distribution and therefore do not produce much noise.

The "alternative variant" (see the end of section 5) is less good here since the direct and caustic lighting is now taken at all points of the volume scattering. Therefore, it is necessary to store in memory all the points of volume scattering on a long path, so that later it is possible to take a convolution of the phase function with direct and caustic lighting from photon maps in these points. If you first trace the rays from the camera (and only then ray trace from the light sources), forming a backward photon map, and the rays from the source are not stored and processed on the fly, this will require unacceptably large memory.

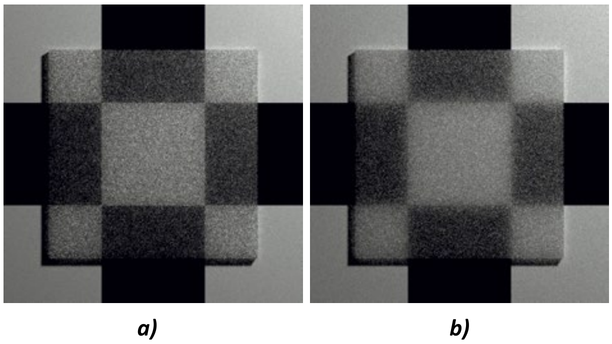

Fig. 2. Camera images for the scene of a turbid medium plate laid upon the paper sheet with a chessboard-like texture. The left panel (a) was calculated in the "standard mode" and the right one (b) was calculated when the volumetric scattering is treated as quasi-specular (in both cases *BDD*=1)

7. RESULTS

7.1. Surface Scattering

Calculations for the modified Cornell Box scene described in section 3 were performed in the quasi-specular mode for *BDD* = 0 and under the same other conditions and for the same time (200 s) as for the standard model. The results for the standard and the quasi-specular methods are shown in Fig. 1.

When applying the proposed method, it was possible to achieve the same low noise of the Cornell Box walls as for the standard method with *BDD* = 0. At the same time, at the bottom of the parallelepipeds, we were able to reduce the noise in about the same way as for the standard method with *BDD* = 1. Thus, in a sense, the best result is obtained.

7.2. For Volume Scattering

The model scene is a 3 mm thick plate laid on a piece of paper with a checkerboard texture, illuminated by a self-luminous sphere above it. The scattering medium inside the plate has a refractive index of 1.5, a scattering coefficient of 7.5 mm^{-1} and a Henyey-Greenstein phase function [13] with parameter $g = 0.9$.

Images calculated the equal time and with the same other parameters are shown in Fig. 2. When using the standard method, the borders of the texture squares are sharp, while they should be blurred. When using the quasi-specular representation of the phase function, these borders are blurred, which corresponds to what one sees.

The left panel (Fig. 2 (a)) was calculated in the "standard model" and the right one (Fig. 2 (b)) was calculated when the volume scattering is treated as quasi-specular. In both cases *BDD* = 1.

REFERENCES

1. H. W. Jensen, Global illumination using photon maps, in Proceedings of the Eurographics Workshop on Rendering Techniques '96, London, UK, Springer-Verlag, 1996. pp. 21–30.

2. H. W. Jensen and P. Christensen, High quality rendering using ray tracing and photon mapping, in ACM SIGGRAPH 2007 Courses, SIGGRAPH '07, (New York, NY, USA), ACM, 2007.

3. T. Hachisuka and H.W. Jensen, Stochastic progressive photon mapping, in ACM SIGGRAPH Asia 2009 Papers, SIGGRAPH Asia '09, New York, NY, USA, ACM, 2009. pp. 141:1–141:8.

4. M. Pharr and G. Humphreys, Physically Based Rendering, Second Edition: From Theory To Implementation. San Francisco, CA, USA: Morgan Kaufmann Publishers Inc., 2nd ed., 2010.

5. Д.Д. Жданов, С.В. Ершов, А.Г. Волобой, Адаптивный выбор глубины трассировки обратного луча в методе двунаправленной стохастической трассировки лучей, in Proceedings of 25-th International Conference on Computer Graphics and Vision, September 22–25, Protvino, 2015. pp. 44–49.

6. J. Vorba, Bidirectional photon mapping, in Proceedings of CESCG 2011: The 15th CentralEuropean Seminar on Computer Graphics, Charles University, Prague, 2011.

7. S. V. Ershov, D.D. Zhdanov, and A.G. Voloboy, Estimation of noise in calculation of scattering medium luminance by MCRT, Mathematica Montisnigri, V14, 2019. pp. 60–73.

8. I. Georgiev, J. Krivanek, T. Davidovic, and P. Slusallek, Light transport simulation with vertex connection and merging, ACM Trans. Graph., Nov. 2012. V31, pp. 192:1–192:10.

9. S. Popov, R. Ramamoorthi, F. Durand, and G. Drettakis, Probabilistic Connections for Bidirectional Path Tracing, Computer Graphics Forum, 2015.

10. N. Dodik, Implementing probabilistic connections for bidirectional path tracing in the Mitsuba Renderer, Sept. 2017.

11. J. T. Kajiya, The rendering equation, in Proceedings of the 13th Annual Conference on Computer Graphics and Interactive Techniques, SIGGRAPH '86, New York, NY, USA, ACM, 1986. pp. 143–150.

12. S. V. Ershov, D.D. Zhdanov, and A.G. Voloboy, Treating diffuse elements as quasi-specular to reduce noise in bi-directional ray tracing, Preprint of the Keldysh Institute of Applied Mathematics, 2018. #122.

13. G. Pomraning, On the Henyey-Greenstein approximation to scattering phase functions, Journal of Quantitative Spectroscopy and Radiative Transfer, 1988. V39, #2, pp. 109–113.

Sergei V. Ershov,
Ph.D. of physical and mathematical sciences. He graduated from the Physics Department of Lomonosov Moscow State University in 1988. He defended his thesis at Keldysh Institute of Applied Mathematics RAS in 1991. Senior Researcher at the Department of Computer graphics and computational optics of Russian Academy of Sciences. Research interests: computer graphics, computational optics, numerical methods of mathematical physics, Monte Carlo methods, ray tracing, diffraction tasks

Dmitry D. Zhdanov,
Ph.D. of physical and mathematical sciences. He graduated from ITMO University in 1984. He defended his thesis at Keldysh Institute of Applied Mathematics RAS in 2006. Currently, he is an Associate Professor at ITMO. Research interests: computer graphics, computational optics, virtual and augmented reality

Alexei G. Voloboy,
Dr. of Phys.-Math. Sciences. He graduated from the Department of mechanics and mathematics of Lomonosov Moscow State University in 1988. He defended his Ph.D. thesis at Keldysh Institute of Applied Mathematics RASin 2005, and his doctoral dissertation in 2012. Leading researcher of the Department of Computer graphics and computational optics of Keldysh Institute of Applied Mathematics RAS. Research interests: computer graphics, computational optics, ray tracing, illuminance simulation

Nikolay B. Deryabin
graduated from the faculty of mechanics and mathematics of Lomonosov Moscow State University in 1977. Research associate of the Department of Computer graphics and computational optics of Keldysh Institute of Applied Mathematics RAS. Research interests: computer graphics, graphic drivers