Treating diffuse elements as quasi-specular to reduce noise in bi-directional ray tracing

S.V. Ershov[1], D.D. Zhdanov[2], A.G. Voloboy[1]
measure@spp.keldysh.ru|ddzhdanov@mail.ru|voloboy@gin.keldysh.ru
[1]Keldysh Institute of Applied Mathematics, Russian Academy of Sciences, 125047 Miusskaya sq. 4, Moscow, Russia;
[2]St. Petersburg National Research University of Information Technologies, Mechanics, and Optics, St. Petersburg, Russia

In bi-directional Monte-Carlo ray tracing using photon maps the backward ray collects luminance when hitting a diffuse surface, and the number of diffuse events is limited by the so-called "backward diffuse depth", BDD. Its value strongly affects the resulting noise. The paper describes a method to choose BDD to reduce noise. The idea is that some surfaces while diffuse are processed as specular, not incrementing the diffuse event counter. Also not all illumination photons are collected at the hit points, and this is also affected by the proposed method. We show how this treatment allows to reduce noise both for surface and volumetric scattering.

Keywords: *bi-directional ray tracing, noise reduction, BDF.*

Abbreviations

MCRT = Monte-Carlo ray tracing

FMCRT = forward Monte-Carlo ray tracing.

BMCRT = backward Monte-Carlo ray tracing.

BDF = bi-directional scattering function. It describes surface luminance as a function of the illumination and observation direction

BDD = backward diffuse depth. It is a specific parameter of a hybrid ray tracing, when FMCRT calculates illumination and BMCRT is used to convert it to the observed luminance. In this method the backward ray usually has a limited "length" and terminates after BDD diffuse events.

1. Introduction

A powerful method of calculation of a virtual camera image is a bi-directional Monte-Carlo ray tracing using photon maps [4]. FMCRT creates photon maps that allow to estimate illumination of scene surfaces and then BMCRT converts them to the camera image. It traces rays from camera and wherever they hit a diffuse surface, they take the local illumination from the photon map, "convolve" it with the local BDF and add the resulting luminance to the image pixel. There are many modifications of this method [3], [8], [6], [1] which all follow the above general idea.

An important parameter of such an approach is the number of operators ("ray events") in that product series that are handled with BMCRT. Usually only diffuse events are counted, and their number is named "backward diffuse depth" (BDD). Efficiency of the approach, i.e. the rate of convergence (or noise level, which is more or less the same) strongly depends on that BDD, and its optimal value is specific for each scene.

A better approach is to let it be different across the scene [1] and even mix calculations with different BDD [1], and frequently it is possible to find it, automatically or manually, so that the calculations are quite efficient. But this is not always and in some cases changing BDD does not help and whichever value we choose the image is highly noisy. Here we also prove why an "adaptive choice" of BDD individually for each ray may result in a biased estimates (i.e. the calculated luminance converges to a wrong value).

We propose a simple method which helps in such cases. The idea of the method is that the scene objects (surface or volumetric) which have some diffuse scattering are separated into two sets: "genuine diffuse" and "quasi-specular". The latter are *usually* those with a narrow, nearly-specular scattering function, but formally the criterion can be *arbitrary*, e.g. one can treat a Lambert surface as quasi-specular. Usually the separation treats each scattering function (BDF for a surface or phase function for a medium) as a *sum* of quasi-specular and genuine diffuse components so that both are not zero.

A quasi-specular scattering does not increase the "event counter" (which terminates the ray when it exceeds the BDD). Also, it is necessary that quasi-specular component of BDF does not convolve with the *indirect* illumination (i.e. FMCRT rays that underwent a genuine diffuse scattering). This change to the BMCRT can be applied to both surface and volumetric scattering to reduce image noise and the memory used.

2. Global illumination equation

Light distribution in a scene is described by a "self-consistent field" equation which can be written in different forms that are named "global illumination equation", "rendering equation" [5] etc.: there is some light field in the scene; it illuminates its surfaces which scatter it. This transformation of the illumination I (light that falls onto a surface) into luminance L (light that goes away from a surface point $\boldsymbol{x}$ in direction $\boldsymbol{v}$) is described by the surface BDF f:

$$L(\boldsymbol{v};\boldsymbol{x}) = \int f(\boldsymbol{x};\boldsymbol{v},\boldsymbol{v}')I(\boldsymbol{v}',\boldsymbol{x})d^2\boldsymbol{v}'$$

Then this light, emitted *from* the surface, propagates over the scene and eventually illuminates its surfaces.

The publication of the conference proceedings was supported by the RFBR, grant No. 18-07-20045\18

This transformation of the luminance (light going away off a surface) into illuminance (light incident onto another surface point) is described by the transport operator. The explicit form of the scattering and transfer operators is inessential here so we use a compact notations:

$$L = \hat{F}I \quad (1)$$
$$I^{(i)} = \hat{T}L \quad (2)$$

Notice that here $\hat{F}$ includes *only diffuse* scattering and thus the "direct" component also contains light that underwent pure specular transformation.

Notice L is **the luminance of the scene surfaces**; the luminance of the **camera image** would be $\mathcal{L} = \hat{\mathcal{S}}L$ where $\hat{\mathcal{S}}$ describes the *pure* specular transformation between the camera and a scene surface (frequently $\hat{\mathcal{S}} = 1$). Here and below we shall only calculate the L. Its transformation into the camera image can be applied after that if needed.

The full illumination consists of indirect and direct components:

$$I = I^{(i)} + I^{(0)} \quad (3)$$

Combining the three above equations we arrive at the form of "rendering equation" by Kajiya [5] we shall use:

$$I = \hat{T}\hat{F}I + I^{(0)} \quad (4)$$

3. Iterative solution (Neumann series) of (4) and BMCRT

A widely used approach in computational optics is a combination of forward and backward ray tracing, when the forward part calculates illumination of the diffuse surfaces I and stores it as, say, photon maps [4], and then the backward part "converts" it in the camera image. Notice that the actual illumination we operate is noisy and this noise is then "transferred" to the image, and its final amplitude strongly depends on how the BMCRT part works.

The simplest way is that we trace rays from camera, terminating then at the first diffuse surface where the surface luminance under the full illumination (= $\hat{F}I$) is calculated and *added* to the pixel luminance. This provides estimate of the surface luminance (1), though the result is not perfect because the estimation of I from FMCRT is usually subjected to (spatial) filtering to reduce noise. Thus all fine illumination details such as highlights are usually lost.

Instead one can first apply the N-th iteration of (4) to the illumination, which leads to

$$L = (\hat{F}\hat{T})^N \hat{F}I + \sum_{k=0}^{N-1} (\hat{F}\hat{T})^k \hat{F}I^{(0)} \quad (5)$$

For the exact illumination field I this gives exactly the same result as the above simple $\hat{F}I$, but for the actual, noisy I the second form is frequently *better* because result in the lower level of the *image noise* thanks to the convolution with a power of $(\hat{F}\hat{T})$. The value of N is termed "**backward diffuse depth**", or **BDD** for short.

The term $\hat{F}I^{(0)}$ is the surface luminance under only the *direct* (including caustic) illumination. The indirect (diffusively scattered) illumination I is accounted only in the last hit point as $(\hat{F}\hat{T})^N\hat{F}I$.

The integral operators can be calculated using Monte-Carlo: we fire rays from camera, they hit a surface, are scattered (left operator $\hat{F}$), propagate the scene (operator $\hat{T}$; notice in the *backward* ray tracing the events counted from camera correspond to the operators *left to right*) and so on, until reach the N-th diffuse surface and terminate. In the k- th diffuse hit point (i.e. just *before* the k-th diffuse scattering) we calculate this surface luminance under the direct (and caustic) illumination ($\hat{F}I^{(0)}$) if $k < N$ and full illumination ($\hat{F}I$) if $k = N$. Then scale the result by the luminance transmission factor due to the specular transformation in $\hat{T}$, and *add* to the pixel luminance. The average over ensemble of camera ray converges to L.

4. How the BDD affects the noise level

Let us first consider a simple scene consisting of a diffuse plate illuminated by a cone light (which emits downwards) enclosed in a matte (diffuse) spherical shade. The camera sees only the plate while the lamp shade is out of the view area. If the lamp shade is small enough we can neglect interreflections.

Illumination of the plate is then entirely *indirect* (light from the diffuse sphere). In case of BDD=0 camera ray collects that illumination in the hit point. There is some moderate noise from the FMCRT (= from photon map), but no more, while for BDD=1 the situation becomes much worse.

Indeed, now in the first hit point (i.e. the plate) we take only the direct (and caustic) illumination which in this model is exactly 0. So *all* contribution to pixel luminance comes only from the *second* diffuse hit, which is only possible when the scattered camera ray hits the lamp shade. But this probability is very low, because the BDF (and thus the scattered light cone) is *wide* while the lamp shade is *small*. Therefore, it is very rare that camera ray brings any luminance to the image pixel, and this means high noise, see [2].

The case of BDD>1 is not better because it also requires camera ray hitting the lamp shade.

Therefore, in our model scene the best is BDD=0 while BDD>0 results in a *very* high noise.

Now let us replace the small lamp shade with something very large. The scene becomes the diffuse plate in the centre of the floor of a large room box with diffuse walls and ceiling illuminated by a grid of cone lights (emitting towards the nearest surface and not towards the room centre), so their light can not reach the plate directly. The floor outside the plate is black. The camera is inside the room and sees only the plate. The whole scene is as the one from Figure 1 *only with-*

Работа опубликована при финансовой поддержке РФФИ, грант 18-07-20045\18

out the light in lampshade. At last, unlike the previous model scene, the bottom plate has now a *sharp* (nearly specular BDF).

Like in the previous example, the plate's illumination is purely *indirect*. The light from the grid of cone emitters undergoes many inter-reflections and so creates a rather uniform illumination of the walls and ceiling, though it is not exactly constant and the areas near the light emitters are brighter.

In case BDD=0 we collect the indirect illumination (from FMCRT!) at the first hit point. Illumination comes from the large box which is illuminated rather uniformly, and thus makes a wide cone. Meanwhile, the BDFs is sharp, so it effectively senses only illumination in direction *close to the mirror reflection of the view ray*. Most of FMCRT rays are be effectively excluded and only few of them will contribute to the pixel luminance which leads to a strong noise.

In case BDD=1 the situation improves. At the first camera hit we collect only the direct (and caustic) illumination, but it is 0 there. After diffuse reflection by the plate the camera ray hits the box and there it takes the full (in fact, again indirect) illumination. Most of camera rays and most of FMCRT rays contribute here, so the noise is low, see the images in [2].

The larger values of BDD work *less good*. Since illumination is *indirect*, the contribution to pixel luminance comes only from the *last* hit point. It can be in the box or in the bottom plate with more or less close probabilities. When it is in the box, it is like for BDD=1 i.e. all works good. But when it hits a bottom plate, all is like for BDD=0 i.e. we take a wide cone illumination of a narrow BDF which results in strong noise.

5. When all ``is worse". Introducing the quasi-specular method

The model scenes from Section 4 admit an optimal BDD for which the noise is low. But there are scenes in which whatever BDD we choose the image is highly noisy.

An example is a *combination* of the two model scenes from Section 4: the bottom plate now has BDF which is a sum of a wide (nearly Lambert) and a sharp (nearly specular component). To see which BDF does which contribution to the image, the Lambert component is pure red and the sharp component is nearly pure green;[1] their integral reflectance is 50% and 40%, respectively. Indirect illumination comes both the room walls and ceiling (illuminated by an array of cone lights), and also by a small "central" lamp shade with a cone light inside, see Figure 1.

Unlike the first scene from 4 the light from the lamp shade also results in interreflections, but when the bottom plate is small (recall the light which hits the room floor outside it is absorbed), it is inessential.

[1] Its normalized color is $(0.1, 1, 0.1)$

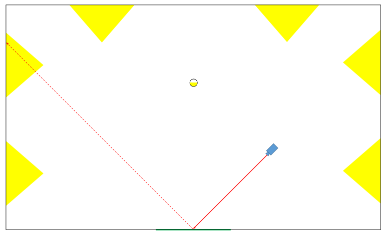

Figure 1. Combined model scene: indirect illumination both from the lamp shade and the room walls and ceiling. This is the side view in correct scale. The yellow cones show emission of lights; the camera view ray and its specular reflection are drawn with red arrows. The circle around the central light is the diffuse lamp shade; the yellow cone inside it mark its light emission (downwards). The thick green strip is the bottom plate.

Repeating the analysis from Section 4 one finds that BDD=0 is bad because of high noise from illumination of a narrow BDF part from the box and BDD=1 is bad because of high noise from illumination of a Lambert BDF from a small lamp shade, see Figure 2. Larger BDDs are not better.

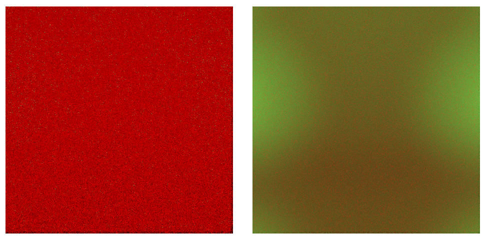

Figure 2. Camera images for scene from Figure 1, calculated during the same time (3000 sec) for BDD=0 (left) and BDD=1 (right). For the area of 100x100 pixels around the centre, the average RGB color is: $(150, 27, 2.71)$ and $(150, 31.6, 3)$. The noise level (relative to the photometric luminance) is 1000%, and 200% respectively.

This, however, does not work [1] which has been proved in [2]. That is, just making the near-specular and off-specular rays to collect illumination according to the ray's BDD produces wrong luminance, and one must apply a more sophisticated selection of the illumination components collected in each hit point.

The key idea of the approach though remains the same: the nearly specular narrow BDF is treated specially, not as a "genuine diffuse"; but in a sense closer to specular ones. Such narrow BDFs (and the whole method of their treatment in BMCRT) are thus named "quasi-specular". The result of calculation of the same model scene with this method is shown in the rightmost panel of Figure 2, see Section 9.1 for more explanations.

The publication of the conference proceedings was supported by the RFBR, grant No. 18-07-20045\18

6. Operator series in presence of quasi-specular BDFs

Now let us come to the formal derivation of what to do with a backward ray when the diffuse BDF is subdivided into the “genuine diffuse” and “quasi-specular” components:

$$\hat{F} = \hat{F}_d + \hat{F}_{qs}$$

The separation can be *arbitrary* (although some separations are advantageous as concerning the noise level and some are not), i.e. the calculated luminance all the same converges to the exact value.

Illumination is also subdivided into three components: **direct** (that was not scattered at all or scattered by pure specular surfaces), **quasi-caustic** (scattered at least once by a specular BDF and never by a diffuse BDF) and **indirect** (scattered at least once by a diffuse BDF and any times by a specular BDF or quasi-specular), i.e.

$$I = I^{(0)} + I^{(qc)} + I^{(i)}$$

Luminance of a surface point $\boldsymbol{x}$ is then:

$$L(\boldsymbol{v}, \boldsymbol{x}) = (\hat{F}_d I)(\boldsymbol{v}, \boldsymbol{x}) + (\hat{F}_{qs} I)(\boldsymbol{v}, \boldsymbol{x})$$

The above luminance expression can be written as

$$L = \hat{F} I = \hat{F}_{qs}\left(I^{(0)} + I^{(qc)}\right) + \hat{F}_{qs} I^{(i)} + \hat{F}_d I \qquad (6)$$

Substituting our decomposition of I into the global illumination equation (4) one arrives at

$$I^{(qc)} + I^{(i)} = \hat{T}\hat{F}_{qs}\left(I^{(0)} + I^{(qc)}\right) + \hat{T}\hat{F}_{qs} I^{(i)} + \hat{T}\hat{F}_d I$$

In the r.h.s., the term $\hat{T}\hat{F}_{qs} I^{(i)} + \hat{T}\hat{F}_d I$ comprises light that underwent at least one *diffuse* scattering while the term $\hat{T}\hat{F}_{qs}\left(I^{(0)} + I^{(qc)}\right)$ comprises light that underwent only specular scattering. In view of the decomposition into diffuse and quasi-caustic illumination, this means

$$I^{(i)} = \hat{T}\left(\hat{F}_d I + \hat{F}_{qs} I^{(i)}\right) \qquad (7)$$

$$I^{(qc)} = \hat{T}\left(\hat{F}_{qs} I^{(qc)} + \hat{F}_{qs} I^{(0)}\right) \qquad (8)$$

from what it follows that

$$I^{(i)} = \left(1 - \hat{T}\hat{F}_{qs}\right)^{-1} \hat{T}\hat{F}_d I \qquad (9)$$

$$I^{(qc)} = \left(1 - \hat{T}\hat{F}_{qs}\right)^{-1} \hat{T}\hat{F}_{qs} I^{(0)} \qquad (10)$$

We assume that if the backward ray underwent scattering by the quasi-specular component, this does *not* increment the diffuse event counter so the ray does not terminate and shall derive which illumination components must taken at which hit points so as “to be compatible” with the above behavior, i.e. so that the mathematical expectation of the BMCRT luminance to match the exact value.

Combining Eq. (6) with Eq. (9) and Eq. (10) we after some tedious yet trivial transformations see that the surface luminance from BMCRT with BDD=N is

$$\begin{aligned} L \;=\; & \hat{F}_{qs}\left(I^{(0)} + I^{(qc)}\right) \\ & + \sum_{k=0}^{N-1}\left(\left(1-\hat{Q}\right)^{-1}\hat{F}_d\hat{T}\right)^k \left(1-\hat{Q}\right)^{-1} \\ & \quad \times \hat{F}_d (I^{(0)} + I^{(qc)}) \\ & + \left(\left(1-\hat{Q}\right)^{-1}\hat{F}_d\hat{T}\right)^N \left(1-\hat{Q}\right)^{-1}\hat{F}_d I \end{aligned} \qquad (11)$$

where

$$\hat{Q} \equiv \hat{F}_{qs}\hat{T} \qquad (12)$$

There is also an **alternative** form, which gives exactly the same result for the *exact* illumination while for a real noisy illumination may give different (in noise level) result:

$$\begin{aligned} L \;=\; & \left(1-\hat{Q}\right)^{-1}\hat{F}_{qs} I^{(0)} \\ & + \sum_{k=0}^{N-1}\left(\left(1-\hat{Q}\right)^{-1}\hat{F}_d\hat{T}\right)^k \left(1-\hat{Q}\right)^{-1} \\ & \quad \times \hat{F}_d (I^{(0)} + I^{(qc)}) \\ & + \left(\left(1-\hat{Q}\right)^{-1}\hat{F}_d\hat{T}\right)^N \left(1-\hat{Q}\right)^{-1}\hat{F}_d I \end{aligned} \qquad (13)$$

Detailed derivation can be found in [2].

7. Integration by paths for BMCRT

Expanding $(1 - \hat{F}_{qs}\hat{T})^{-1}$ in Eqs. (11) via Neumann series we see that e.g. for BDD=2

$$\begin{aligned} L \;=\; & \hat{F}(I^{(0)} + I^{(qc)}) \\ & + \sum_{k=1}^{\infty} \hat{Q}^k \hat{F}_d (I^{(0)} + I^{(qc)}) \\ & + \sum_{k,m=0}^{\infty} \hat{Q}^k \hat{F}_d \hat{T} \hat{Q}^m \hat{F}_d (I^{(0)} + I^{(qc)}) \\ & + \sum_{k,m,n=0}^{\infty} \hat{Q}^k \hat{F}_d \hat{T} \hat{Q}^m \hat{F}_d \hat{T} \hat{Q}^n \hat{F}_d I \end{aligned} \qquad (14)$$

In this expression a term like $\hat{Q}^k \hat{F}_d \hat{T} \hat{Q} \hat{F}_d (I^{(0)} + I^{(qc)})$ means that:

1. The ***final*** transformation (“before camera”) light is subjected to $k \geq 0$ quasi-specular transformations with corresponding transport term, i.e. $\hat{Q}^k = (\hat{F}_{qs}\hat{T})^k$
2. ***Before*** that (i.e. “further from camera”) light is subjected to the *diffuse* transformation with corresponding transport term, i.e. $\hat{F}_d\hat{T}$
3. ***Before*** that (i.e. more “further from camera”) are light is subjected to $m \geq 0$ quasi-specular transformations with corresponding transport term, i.e. $\hat{Q}^m = (\hat{F}_{qs}\hat{T})^m$
4. And all the above is applied to the convolution of the direct and quasi-caustic illumination with the *diffuse* BDF i.e. to $\hat{F}_d\left(I^{(0)} + I^{(qc)}\right)$.

The action of the integral operators $\hat{F}_d$ and $\hat{F}_{qs}$ can be estimated with the BMCRT, tracing ray from camera. Then the first ray transformation events correspond

Работа опубликована при финансовой поддержке РФФИ, грант 18-07-20045\18

to the *leftmost* operators in a product and the last ray events — to the *rightmost* operators. That is, our term $\hat{Q}^k \hat{F}_d \hat{T} \hat{Q} \hat{F}_d (I^{(0)} + I^{(qc)})$ is estimated from the rays that **first** underwent $k \geq 0$ quasi-specular events, **then** one diffuse event, then $m \geq 0$ quasi-specular events and **after** that take the luminance of the *diffuse* BDF under direct and quasi-caustic illumination (i.e. the diffuse illumination part is ignored).

Detailed derivation can be found in [2].

The other terms in Eq. (14) can be interpreted and then estimated with BMCRT similarly. This leads to the following algorithm of processing BMCRT rays:

- before (and including!) the first not pure specular event we take $\hat{F}\left(I^{(0)} + I^{(qc)}\right)$;
- after the first quasi-specular event and up to the 2nd diffuse event we take $\hat{F}_d\left(I^{(0)} + I^{(qc)}\right)$;
- *after* the 2nd diffuse event we take $\hat{F}_d I$
- at the 3rd diffuse event we stop

The case of another BDD, as well as the alternative form i.e. Eq. (11) are considered similarly. When applying BMCRT, we trace rays from camera so that they undergo BDD+1 diffuse events (unless are absorbed prematurely) and then stop. Whenever a ray hits a surface which has some diffuse or quasi-specular BDF, the ray takes the convolution of *some part* of the BDF with *some part* of illumination, as summarized in the Table 1:

Table 1. In which hit points and for which illumination component the "genuine diffuse" BDF part $\hat{F}_d$ is used instead of the full $\hat{F}$. "QS" mean "quasi-specular".

	Main variant	**Alternative variant**
Direct, caustic	After the 1st QS event	After the 1st diffuse event
Quasi-caustic		Always
Indirect	Always	

8. Volumetric scattering

8.1 Standard method

Here all works like for surface scattering. Suppose for example that the camera is inside a turbid medium and BDD=1. Then camera ray propagates in the medium and when it undergoes the first volumetric scattering, it takes *direct* and *caustic* illumination and "convolves" it with the phase function. In the second volumetric scattering it takes *full* illumination and "convolves" it with the phase function.

Let us consider a model scene in which camera looks through a layer of turbid medium onto some illuminated object. Let also the turbid medium is absorption-free and has a large scattering coefficient, so that camera ray undergoes many (volumetric) scattering events before it penetrates the medium and reaches the object. Let, at last, phase function be sharp enough so that each scattering changes the ray direction only slightly.

For the "standard" method with a small BDD, the camera ray terminates close to the ray origin (i.e. camera), and the ray does not reach the object. Also, since the phase function is narrow, its convolution with that illumination which came from the scene surfaces results in strong noise, cf. Section 4.

We must set a *large* BDD so that camera ray to *leave* the medium, reach a scene surface behind and also not to collect indirect illumination in the volumetric scene points (which due to sharp phase function creates strong noise). But then we must keep in memory all the volumetric scattering points (because they still collect direct and caustic illumination), and this is usually too expensive.

8.2 "Quasi-specular medium"

Applying the quasi-specular approach to the volumetric scattering greatly improves the situation with the above scene.

Suppose that the scene *surfaces* are not quasi-specular. Let also the *whole* phase function be treated as quasi-specular, i.e. its "genuine diffuse" part $\hat{F}_d = 0$. Then, until the camera ray travels inside the medium, it undergoes only quasi-specular events and thus does not increment the "diffuse counter". As a result it penetrates the medium layer and reaches for a scene surface.

Indirect illumination now is the light after diffuse scattering by scene surfaces; *quasi-caustic* is the light which underwent at least one volumetric scattering (and any specular events, but no diffuse surface scattering).

Since $\hat{F}_d = 0$, the "main variant" from Table 1 implies that:

- the *indirect* illumination (=from scene surfaces) is effectively *ignored* inside the medium
- the *direct, caustic* and *quasi-caustic* illumination is taken only up to the *first* volumetric scattering, i.e. in fact for the *first* volumetric event only.

As a result, for *any* BDD the camera ray reaches the scene surfaces behind the medium layer; only *one* volumetric ray event (the first one) is remembered; there is no noisy convolution of the sharp phase function with the indirect illumination from scene surfaces.

There is, though, the convolution of phase function with the direct, caustic and scattered by the medium light. But they all do not have a *wide* angular distribution that would result in strong noise.

The "alternative" variant from Table 1 is worse because the direct and caustic illumination *will be* taken "up to the first *diffuse* event" and since the volumetric scattering is quasi-specular, that happens only *after* the camera rays leaves the medium. We must therefore

The publication of the conference proceedings was supported by the RFBR, grant No. 18-07-20045\18

store *all* the ray scattering events inside the medium to take direct and caustic illumination in that points which is expensive in memory.

9. Results

9.1 Surface case

The model scene is from Figure 1. Calculation for the quasi-specular mode were performed for BDD=0 and the same other conditions as those for the “standard mode”. The results for the standard mode are in Figure 2 and for the quasi-specular mode in Figure 3.

One can see the noise level is about threefold lower than for the best case without the quasi-specular method. And, which is more essential, the calculated image has a better estimation of color: for BDD=0 the image looks mainly red with rare bright green dots; for BDD=1 it looks mainly green with rare bright red dots, and it is only with quasi-specular method that we see a “mixture” of red and green. In the meanwhile the *averages* obtained in the three simulations are about the same, which means that the result e.g. for the “standard BDD=1” just has very bright red dots which “on average” would give the correct value. That is, in the standard method achieves the correct average by very rare very bright peaks, i.e. the worst sort of noise.

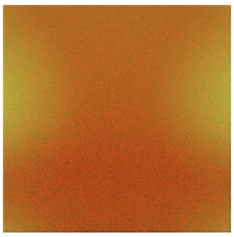

Figure 3. Camera image for scene from Figure 1, calculated for quasi-specular method with BDD=0. For the area of 100x100 pixels around the centre, the average RGB color is (150, 31.4, 3). The noise level (relative to the photometric luminance) is 76%.

9.2 Turbid medium

The model scene consisted of a plane parallel plate of thickness 3 mm laid upon a paper sheet with a chessboard-like texture, illuminated by a self-emitting sphere above it. The plate’s medium has refraction index 1.5 and scattering coefficient 7.5 mm^{-1}. The phase function is the Henyey-Greenstein one [7] with $g = 0.9$. The images calculated during the same time (and other settings) are presented in Figure 4; one can see that in the “standard mode” the boundaries of the texture squares seen through the plate are *sharp* while in reality they must blur. In the image calculated with the quasi-specular representation of the phase function these boundaries are smoothed (as they should be).

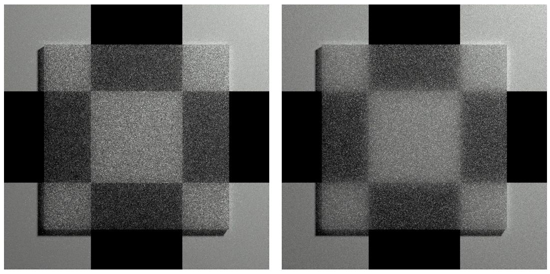

Figure 4. Camera images for the scene with a plate of turbid medium laid upon a chessboard-like texture. The left image is the “standard mode” of volumetric scattering, when it is treated as diffuse. The right panel is when the volumetric scattering is purely quasi-specular. In both cases BDD=1.

Acknowledgments

The work was partially supported by RFBR, Grant 18-01-00569.

10. References

[1] 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, 2015.*, pages 44–49, Protvino, 2015.

[2] S.V. Ershov, D.D. Zhdanov, A.G. Voloboy, M.I. Sorokin. Treating diffuse elements as quasi-specular to reduce noise in bi-directional ray tracing. *Preprint of the Keldysh Institute of Applied Mathematics*, No. 122, doi:10.20948/prepr-2018-122-e, 2018.

[3] P. Christensen, H.W. Jensen. High quality rendering using ray tracing and photon mapping. In *ACM SIGGRAPH 2007 Courses*, SIGGRAPH ’07, New York, NY, USA, 2007. ACM.

[4] H.W. Jensen. Global illumination using photon maps. In *Proceedings of the Eurographics Workshop on Rendering Techniques ’96*, pages 21–30, London, UK, UK, 1996. Springer-Verlag.

[5] J.T. Kajiya. The rendering equation. In *Proceedings of the 13th Annual Conference on Computer Graphics and Interactive Techniques*, SIGGRAPH ’86, pages 143–150, New York, NY, USA, 1986. ACM.

[6] G. Humphreys, M. Pharr. *Physically Based Rendering, Second Edition: From Theory To Implementation.* Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2nd edition, 2010.

[7] G.C. Pomraning. On the Henyey-Greenstein approximation to scattering phase functions. *Journal of Quantitative Spectroscopy and Radiative Transfer*, 39(2):109 – 113, 1988.

[8] H.W. Jensen T. Hachisuka. Stochastic progressive photon mapping. In *ACM SIGGRAPH Asia 2009 Papers*, SIGGRAPH Asia ’09, pages 141:1–141:8, New York, NY, USA, 2009. ACM.

Работа опубликована при финансовой поддержке РФФИ, грант 18-07-20045\18