

A ZERO-VARIANCE-BASED SAMPLING SCHEME FOR MONTE CARLO SUBSURFACE SCATTERING


Jaroslav Křivánek
Charles University in Prague

Eugene d'Eon
Weta Digital

The work was done while both authors were with **Weta Digital**



Monte Carlo subsurface scattering



Křivánek and d'Eon - Zero-variance-based MC SSS


2

- This work was started during my last year's stay at Weta Digital and deals with the calculation of subsurface scattering using Monte Carlo simulation as opposed to an analytic BSSRDF model.

Motivation


- **Classical random walk**

- Oblivious to the boundary



- **Goal**

- Guide paths toward the boundary



- The MC simulation involves a random walk, where the light path enters under the surface, scatters around in a way that is totally oblivious to where the surface actually is and then eventually, it may get out and connect to light sources, but often it just dies somewhere under the surface.
- The visualization (top right) shows a superposition of a number of such random walks and you can see that they have the tendency to wander quite far from the entry point.
- We thought this wasn't a particularly smart way of doing things and we wanted to somehow inform the random walk that it should eventually get back to the surface because that's where the light sources can be found.
- The method we end up developing does exactly that. You can see that the resulting random walks are much more concentrated around the entry point, which is what we set to achieve.
- (Note each path is rasterized with the same opacity, regardless of sample weight.)

Teaser

Classical random walk

Oblivious to the boundary



New method

Guides paths toward the boundary



- Variance reduction
- Efficiency improvement (shorter paths on average)

- However, the method we use does more than just guiding paths toward the surface, thus improving efficiency.
- It is mainly a **variance reduction** scheme. The improved efficiency only comes as a nice extra benefit to it.
- You can see one result here, the left image was rendered with a classic MC random walk, and the right one using our improved subsurface sampling scheme.

PREVIOUS WORK

Previous work in graphics

- Diffusion approximation for SSS [Jensen et al. '01]



- Traditionally, subsurface scattering in graphics has been simulated using analytical BSSRDF models based on the diffusion approximation.
- This approach was introduced to graphics by Jensen et al. in 2001.

Previous work in graphics

- **Diffusion approximation** for SSS

- Multi-layered materials
[Donner and Jensen '05]
- Hybrid MC–diffusion
[Li et al. '05, Donner and Jensen '07]
- Extended source
[d'Eon and Irving '11, Habel et al. '13]
- Direction dipole
[Frisvald et al. '14]



- There's plenty of follow-up work on this topic, some of which is mentioned on this slide.

Diffusion-based SSS – Limitations

- Assumptions do not hold in practice (flat, semi-infinite medium)
- Inaccuracies of the diffusion approximation

- However, the diffusion-based BSSRDF models come with some serious limitations.
- First, the model is usually valid only for semi-infinite media with a flat boundary, and this assumption is rarely met in practice. When rendering a human head, this would be problematic especially around the nose and the ears.
- But even if the assumption held, there are inaccuracies of the diffusion approximation itself: after all, it's an *approximation*.


Monte Carlo SSS

- Accurate, no assumptions about the geometry
- Fits smoothly into a physically-based path tracer
- **BUT** slow
 - Hundreds of scattering events to get an accurate answer
- Similar issues tackled in **neutron transport**

- For these reasons, an explicit Monte Carlo (MC) simulation of the subsurface scattering process is an interesting alternative to the approximate analytical BSSRDFS.
- With this approach, many approximations and assumptions go away and the integration into a physically-based path tracer is very natural.
- But the MC simulation can be very slow because a high number of scattering events have to be simulated for each path. Furthermore, the result is just a random sample with an associated variance (=noise), whereas evaluating a BSSRDF is relatively fast and gives a noise-free result.
- Can we do something to alleviate the limitations of MC sampling?
- As is often the case, it turns out that neutron transport research has investigated a similar problem many years ago.

Previous work in neutron transport

- “Deep penetration problems” (reactor shield design)
 - Blind MC: **One in a billion** particles makes it through



Křivánek and d'Eon - Zero-variance-based MC SSS

10


- The so called “deep penetration problems” used in reactor shield design are particularly closely related.
- By design, only a tiny fraction of the incident radiation is allowed to pass through a reactor shield.
- For example, in a blind MC simulation only one in a billion particles would make it through, which makes the classic MC simulation totally hopeless.

Previous work in neutron transport

■ Path stretching

[Clark '66, Ponti '71, Spanier '71]

- heuristic, manual “stretching parameter” setting



Křivánek and d'Eon - Zero-variance-based MC SSS

11

- One of the approaches to solve this issue is the so-called “path stretching”.
- The idea is to advance the particles toward the outside by artificially stretching the sampled distance whenever the particle in a MC simulation points toward the exterior, and to shrink it when the particle is directed to the interior. To compensate for this, one needs to adjust the weight of the particle, because its behavior no longer follows the laws of physics.
- Path stretching has originally been derived heuristically and relied on an ad-hoc parameter (the ‘strength’ of the stretching). While it often worked great, it could actually deteriorate the result (increase variance) when the stretching parameter wasn’t set judiciously.

Zero-variance random walk theory

- Zero variance random walk theory
[Kahn '54, Kalos and Whitlock '08, Hoogenboom '08, Booth '11]
- [Dwivedi '81]
 - Synergistic **path stretching** and **angular sampling**
 - Solid theory, no heuristics
- We apply these ideas in SSS in computer graphics

- This is where the theory of zero-variance random walks comes into play.
- The idea that a random walk can be constructed in such a way that it always yields the correct answer with absolutely no variance has been around for almost as long as MC methods themselves.
- Despite the zero variance theory being old, Hoogenboom's recent 2008 NSE article (*Zero-variance Monte Carlo Schemes Revisited*) is very important: he corrects some misconceptions about the uniqueness of zero-variance walk construction that have lingered for several decades, and includes discussions about boundary crossing and track-length estimators as well.
- Booth's 2012 article (*Common misconceptions in Monte Carlo particle transport*) further clarifies and generalizes some of the concepts of zero-variance schemes. He argues that Hoogenbooms' conclusion concerning the uniqueness of the zero-variance constructions are not correct and that there are multiple ways in which a zero-variance walk can be constructed.
- Dwivedi [1981] was the first to apply the theory of zero-variance random walks to deep penetration problems.
- He has shown how the heuristic path stretching automatically follows from the theory, while giving a clear answer to the parameter setting.
- He was also the first to show that to robustly reduce variance, the path stretching needs to be combined with an appropriate angular sampling.
- While he conceived his work with reactor shield design problems in mind, we apply his ideas to subsurface light transport and make further improvements.

THEORY OF ZERO-VARIANCE RANDOM WALKS

- Because we believe that the theory of zero-variance random walks is important, we start by exposing it in a general setting, without any reference to our particular problem.

WARNING!


- Zero-variance walk much harder than calculating the desired answer!
 - Requires knowing the solution **everywhere**

- **Approximate zero-variance walks yield low variance**
 - [Dwivedi '81]
 - This work

- Before starting, let me point out that constructing a zero-variance random walk is a hard problem. In fact, it is much harder than calculating the desired answer itself.
- In order to construct a zero variance random walk, one needs to know the desired **solution everywhere** and be able to generate samples proportional to it.
- So at the first sight, this looks like a fairly useless thing to do.
- But, we can achieve practical results – something that approaches the zero-variance ideal – by using suitable approximations.
- And this is what Dwivedi did in neutron transport and what we propose to do in graphics.

Transport equation

$$L(\mathbf{r}) = L^e(\mathbf{r}) + \int L(\mathbf{r}')T(\mathbf{r}' \rightarrow \mathbf{r})d\mathbf{r}'$$




Křivánek and d'Eon - Zero-variance-based MC SSS

15

- Suppose we have a volume with a volumetric light source.
- Particles (neutrons, photons) in this volume are defined by their position and direction of travel that we will, for the sake of conciseness, denote as \mathbf{r} (the “phase-space position”).
- The volume rendering equation can be written in a simple way as shown at the top of the slide. In this more general form, we refer to it simply as the transport equation.
- It states that the radiance at \mathbf{r} (please recall that \mathbf{r} is both a spatial position and a direction) is the self-emission at \mathbf{r} plus the radiance transported to \mathbf{r} from all other positions \mathbf{r}' of the volume.
- The transport kernel T consists of a **spatial** as well as the **angular** (=directional) part. The spatial part describes the effect of light attenuation along the line from \mathbf{r}' to \mathbf{r} , and it is equal to the transmittance between \mathbf{r} and \mathbf{r}' . The angular part described the directional properties of scattering in the medium and is equal to the phase function at \mathbf{r} .

Solving the transport equation

$$L(\mathbf{r}) = L^e(\mathbf{r}) + \int L(\mathbf{r}')T(\mathbf{r}' \rightarrow \mathbf{r})d\mathbf{r}'$$




$$\mathbf{r} = (\mathbf{x}, \omega)$$

- Solving the transport equation using Monte Carlo involves a recursive application of MC integration to evaluate the integral on the right hand side.
- To estimate the radiance at \mathbf{r} , we sample a phase-space position \mathbf{r}' at random, evaluate the integrand and divide by the pdf of choosing \mathbf{r}' . This involves evaluating the transport kernel and also the radiance at \mathbf{r}' . But the radiance at \mathbf{r}' is unknown and we have to again evaluate it using a recursive MC estimation of the integral, this time at \mathbf{r}' .
- This leads to the classic random walk as we know it from the path tracing algorithm.

Zero-variance estimator

$$L(\mathbf{r}) = L^c(\mathbf{r}) + \int L(\mathbf{r}')T(\mathbf{r}' \rightarrow \mathbf{r})d\mathbf{r}'$$



1. Terminate with prob.:

$$P_{\text{term}}(\mathbf{r}) = \frac{L^c(\mathbf{r})}{L(\mathbf{r})}$$

2. Score **only** when terminating

3. Sample \mathbf{r}' with pdf:

$$p(\mathbf{r}' | \mathbf{r}) \propto L(\mathbf{r}')T(\mathbf{r}' \rightarrow \mathbf{r})$$

- Clearly, if the result of the estimator were always the same, it would have zero variance.
- Let us now see how to make this calculation return the same result no matter what random walk we sample.
- What I describe here is just one way to construct a zero variance random walk, but it is not the only one as pointed out by Booth [2011].
- **Rule 1:** The walk termination probability must be equal to the ratio of the self-emission to total radiance. In our words, the walk is only allowed to stop on a light source. The more indirect illumination, the less likely it is that the walk will stop.
- **Rule 2:** The intermediate vertices of the walk must not make any contribution to the result; we only accumulate emission on terminating the walk. From the path integral point of view, one single random walk yields just a single light transport path (as opposed to the usual practice in, say, path tracing).
- **Rule 3:** When continuing the walk, we pick the next phase-space position \mathbf{r}' from the unique pdf that is exactly proportional to the integrand in the transport equation.
- Why is this difficult? One needs to know the answer $L(\mathbf{r})$ to decide the termination. Furthermore, one needs to sample from the product of the correct solution (which, therefore, needs to be known everywhere!!!) and the transport kernel to pick the next phase-space position.

- Let us now **prove** that these rules lead to a zero-variance calculation. Our goal is to calculate $L(\mathbf{r})$ which is given by the sum of two terms. Let us calculate the sum with MC summation by picking only one of the two terms. To achieve zero-variance, the probability of picking the terms must be exactly proportional to the their respective contributions to the sum.

Proof of zero-variance

$$L(\mathbf{r}) = L^e(\mathbf{r}) + \int L(\mathbf{r}')T(\mathbf{r}' \rightarrow \mathbf{r})d\mathbf{r}'$$

- If the $L^e(\mathbf{r})$ term selected

- Return

$$\frac{L^e(\mathbf{r})}{P_{\text{term}}(\mathbf{r})} = \frac{L^e(\mathbf{r})}{L^e(\mathbf{r})/L(\mathbf{r})} = L(\mathbf{r})$$

- Terminate

1. Terminate with prob.:

$$P_{\text{term}}(\mathbf{r}) = \frac{L^e(\mathbf{r})}{L(\mathbf{r})}$$

2. Score **only** when terminating

3. Sample \mathbf{r}' with pdf:

$$p(\mathbf{r}' | \mathbf{r}) \propto L(\mathbf{r}')T(\mathbf{r}' \rightarrow \mathbf{r})$$

- Let's see what happens if we pick the first term.
- The probability of choosing this term must be proportional to the self-emission in order to obtain a zero-variance estimator of the sum.
- Therefore, it is given by the self-emission L^e divided by the sum (so that it is a properly normalized probability).
- Notice that this probability is exactly equal to the termination probability of the zero-variance walk (cf. Rule 1).
- The estimator of the sum, in this case, is given by the summand, L^e , divided by the probability of choosing this first term.
- This fraction is exactly equal to the desired answer $L(\mathbf{r})$ because the self-emission L^e cancels out.
- And because evaluating the first term of the sum only involves nothing but reading the value of L^e from the scene database, there is really nothing left to do, so the walk terminates.
- Note that in this case, the complicated thing is calculating the termination probability, as it requires knowing the correct solution $L(\mathbf{r})$.

Proof of zero-variance

$$L(\mathbf{r}) = L^c(\mathbf{r}) + \int L(\mathbf{r}')T(\mathbf{r}' \rightarrow \mathbf{r})d\mathbf{r}'$$

■ Else:

□ Return

$$\frac{1}{1 - P_{\text{term}}(\mathbf{r})} \left\langle \int L(\mathbf{r}')T(\mathbf{r}' \rightarrow \mathbf{r})d\mathbf{r}' \right\rangle =$$

1. Terminate with prob.:

$$P_{\text{term}}(\mathbf{r}) = \frac{L^c(\mathbf{r})}{L(\mathbf{r})}$$

2. Score **only** when terminating

3. Sample \mathbf{r}' with pdf:

$$p(\mathbf{r}' | \mathbf{r}) \propto L(\mathbf{r}')T(\mathbf{r}' \rightarrow \mathbf{r})$$

- If we choose the second term, the estimator is given by the 2nd term of the sum (the integral) divided by the probability of choosing the 2nd term (which is 1 minus the termination probability).
- We do not know the value of the integral so we replace it by its estimator.
- But in order for the entire random walk to yield a zero-variance result, the estimator of this integral itself must have zero variance.
- This is achieved by sampling the next position \mathbf{r}' from a pdf exactly proportional to the integrand.
- Of course, this pdf must be properly normalized, so it is given by the integrand divided by the integral itself.
- The actual equations for these ideas are provided on the supplemental slide at the end of the slideshow.
- After some manipulations (see the supplemental slide), the estimator in this case ends up being given by the formula: $L(\mathbf{r}) \langle L(\mathbf{r}') \rangle / L(\mathbf{r}')$.
- The estimator $\langle L(\mathbf{r}') \rangle$ is evaluated by recursively repeating the same zero-variance procedure described here for \mathbf{r} .
- Sooner or later in the recursion, the first summand (self-emission) will be chosen, the walk will terminate, and start returning the actual value of $L(\mathbf{r}')$ as a result of the estimator $\langle L(\mathbf{r}') \rangle$.
- As a consequence, $\langle L(\mathbf{r}') \rangle / L(\mathbf{r}')$ will cancel out and we are, again, left with the correct answer $L(\mathbf{r})$.

- This completes the proof that the given three rules yield a random walk that always returns the same correct answer.

Discussion

- Requires knowing $L(\mathbf{r})$ everywhere
- **Approximate $L(\mathbf{r})$ yields low variance**
- Use in graphics
 - MC calculation [Jensen 1995, Vorba et al. 2014]
 - Analytic solution: **this work**

- It should now be clearer why constructing the zero-variance walk is actually that hard: We needed to know the answer $L(\mathbf{r})$ to decide the termination criterion. Furthermore, we needed to sample from the product of the correct answer and the transport kernel to pick the next phase-space position.
- However, by using an approximate solution to perform these two operations, we can achieve a random walk that is close to the zero-variance ideal.
- One way to obtain the approximate solution is through a rough Monte Carlo pre-calculation.
- This idea has been used before in graphics, though without the reference to the theory of zero-variance walks (e.g. [Jensen 1995, Vorba et al. 2014] and many others).
- These works use sampling according to the product of the BRDF and the incoming radiance to choose the direction of the next ray in path tracing.
- This corresponds to choosing the next phase-space position \mathbf{r}' with a pdf proportional to the integrand of the transport equation.
- However, what many of these previous works have missed is that in a truly zero-variance walk, the rules for terminating the walk must also be modified from the usual ones.
- Probably the first work in graphics to consider the path termination issue in conjunction with approximating a zero-variance scheme is that of Vorba et al. [2014].


- Another option for obtaining the approximate solution is using an analytic solution for a canonical configuration for which a closed-form solution exists and which that is similar to the practical configuration under consideration.
- This is the approach that we use for simulating subsurface scattering.

(NEAR)-ZERO-VARIANCE SOLUTION FOR HALF- SPACE PROBLEMS

- Let us now follow Dwivedi's work and approximate a zero variance walk for subsurface scattering.

Setup


■ (Unidirectional) path tracing



- We apply the technique in a unidirectional path tracer.
- The zero-variance-based random walk is used only for the part of the path under the surface.
- The rest of the path is not affected at all and follows the same rules as in a regular path tracer.

Assumptions

- Flat, semi-infinite medium
- Uniform, white-sky illumination




- Since we do not know what the path tracer will encounter after escaping from the surface, we assume, for the sake of constructing the zero-variance walk, white-sky illumination, uniform in the directional and spatial domains.
- This is equivalent to saying that 'escape' from the medium is our only source of importance that guides the random walk.
- Note that this is not a necessary step: the general theory ('Caseology' on the next slide) permits knowing (in theory) the full directional radiance distribution inside some volume due to a particular light source(s) outside of it. In this much more complicated case, the approximate internal radiance solution would guide the subsurface sub-paths not only to try to escape the medium, but also tend to positions which permit angle selections that then leave the medium in a direction that tends to hit the light. However, this would be quite complex in practice so we decided not to pursue this option.

Analytical radiance solution

■ “Caseology”, singular eigenfunctions

[Case 1960, McCormick and Kuscer 1973]

Singular eigenfunctions



$$(x, \mu) = \int_{-1}^1 \underbrace{\phi(\nu, \mu) e^{-x/\nu}}_{\text{transient terms}} d\nu + \underbrace{\phi_0(\nu_0, \mu) e^{-x/\nu_0}}_{\text{asymptotic terms}}$$

- In order to be able to build upon the zero-variance theory in constructing the random walk, we need an approximation of the radiance solution under the surface.
- Note that the Monte Carlo estimator that we construct is unbiased irrespective of the accuracy of this approximate solution; the accuracy only affects the variance reduction but not the unbiasedness of the resulting estimator.
- For this reason, we choose to approximate the true sub-surface distribution of radiance using a solution for a half-space with flat boundary.
- In this setting, the phase space position \mathbf{r} can be specified by the (optical) depth under the surface, x , and the direction cosine ν .
- The exact half-space solutions that we use can be derived in several ways, including what is known as Caseology, named after Kenneth Case who was on the theoretical division at Los Alamos building the first bombs. In 1960 Case studied in detail the 'spectrum' of the transport operator in the plane-parallel case, and derived an expansion into the discrete asymptotic diffusion mode, and the continuous spectrum of 'singular eigenfunctions' - so named because the angular distributions (the phi-functions on the slide) for the transient terms are singular. This study of the structure of exact solutions in transport is incredibly insightful and we are the first, to our knowledge, to exploit it directly in a rendering technique.
- Each eigenfunction of the spectrum has the same form: the spatial term that

only depends on the depth under the surface is a simple exponential, with different eigenfunctions having different decay rates. The angular term then corresponds to integrating the spatial term along a line from a given depth in a given direction.

Approximate solution

- Use only the discrete asymptotic term



$$L(x, \mu) =$$


$$\phi_0(\nu_0, \mu) e^{-x/\nu_0}$$

asymptotic terms

- Since the complete solution is rather complex – being given by an integral over the spectrum – Dwivedi choose to only use the discrete (asymptotic) mode of the spectrum, disregarding its continuous part.
- This is a reasonable assumption because the spatial part of the discrete mode has a slower decay rate than the rest of the spectrum. Therefore, as one moves away from the boundary, this terms dominates all the other terms and *asymptotically* becomes the correct solution.

Approximate solution

- Use only the discrete asymptotic term




- This slide illustrates in a log-plot the accuracy of approximating the spatial part of the sub-surface solution (the dotted lines) with a single exponential (the straight dashed lines).
- We can see that the approximation is very accurate for the low absorption case (high albedo). As the absorption increases (albedo decreases) the approximation becomes less accurate, especially near the boundary, where the transient terms (the continuous part of the Case's spectrum) dominates the solution.

Distance sampling

$$p(s) \propto \underbrace{\exp(-s\sigma)}_{\text{transmittance}} \cdot \underbrace{\exp(-x/v_0)}_{\text{solution}}$$

$$\propto \exp(-s\sigma') \quad \sigma' = \sigma(1 - \mu/v_0)$$


path stretching!



- We follow Dwivedi and use only the discrete, asymptotic term as an approximation of the true solution under the surface for constructing the random walk.
- Given a particle as a location $\mathbf{r} = (x, u)$ under the surface, we want to randomly sample its next position $\mathbf{r}' = (x', u')$.
- This is decomposed into sampling the next collision distance along the line from x in the direction u , followed by sampling the new particle direction.
- When sampling the collision distance, the zero-variance theory tells us that we should do this with a pdf proportional to the product of the spatial part of the transport kernel (i.e. transmittance = exponential) and the solution (which we assume to the Case's discrete mode = exponential).
- This yields an extremely convenient result, where we sample from a simple exponential pdf, where the transport coefficient is modified from its true value using the equation shown on the slide.
- Note that this is exactly equivalent to the idea of path stretching: for directions toward the boundary the sampled distance will be stretched, whereas it will be shortened for directions away from the boundary.
- Note that the v_0 parameter is a constant here, because it only depends on the single scattering albedo of the medium.

Directional distribution

$$p(\mu) \propto \frac{v_0}{v_0 - \mu}$$




- The work of Dwivedi was the first to point out that path stretching must be accompanied by a corresponding modification of the direction sampling in order to achieve a robust variance reduction.
- This follows naturally from the zero-variance theory: we need to sample the direction from the product of the angular part of the transport kernel (=phase function) and the angular part of the true solution.
- In our work, we assume isotropic media with a constant phase function, so we only need to sample from a pdf proportional to the angular part of the solution.
- This can be derived simply by integrating the spatial part of the solution (=exponential) along lines of different direction.
- Furthermore, in deriving the angular solution, Dwivedi assumes that the spatial solution extends beyond the boundary surface, so he does not need to treat the direction toward the boundary as a special case.
- This yields an extremely elegant solution for the angular part, which is independent of the depth under the surface, and can be easily normalized and used for sampling in a closed form.
- Note that because we only assumed 'escape' as our goal, the 'guided' direction selection only involves modifying the selection of the direction cosine μ : the azimuthal angle selections at each step are chosen uniformly. This is one feature that would change if, say, you knew the source of light outside the medium came from a particular direction. In these cases, more advanced deterministic solutions interior to the medium would be required, and works of


Jakob et al. [2014] (*A Comprehensive Framework for Rendering Layered Materials*) and d'Eon [2014] (*A Dual-Beam 3D Searchlight BSSRDF*) could both be used to produce such importance functions for walk guidance (for the plane-parallel directional light source case, or the point source near a subsurface surface case).

Directional distribution

albedo = 0.95




albedo = 0.2



- These plots show the resulting angular pdf (the horizontal axis corresponds to the direction cosine and the vertical axis is the pdf value).
- We can see that for low absorption (high albedo) the solution is mostly uniform, which corresponds to the well-known fact that the diffusive multiple scattering in low-absorption media leads to a solution that is mostly uniform in directions.
- For high absorption, we can see a pronounced peak of the distribution with strong preference to sampling directions toward the boundary.






RESULTS

Semi-infinite half-space test



- Before using the method in rendering, we have tested it in a simple simulator of light transport in a half-space.
- On this slide, we demonstrate the effect that the new method has on the trajectory of light paths under the surface.
- In the classical sampling the paths tend to wander quite far from the point of entry, while the new sampling scheme concentrates most of the sampling effort around the entry point. This effect is more pronounced for higher absorption levels (lower albedos) because in this regime, the diffusive multiple scattering has only low effect on the final result.

Semi-infinite half-space test


	albedo	0.4	0.8	0.95
new scheme (Dwivedi)				
variance		0.020 (1×)	0.034 (2.7×)	0.025 (5.3×)
classical sampling				
variance		0.020	0.093	0.132

J. Křivánek - Zero variance walks for subsurface scattering

32

- The images on this slide correspond to rendering subsurface scattering on the flat, index-matched boundary of a semi-infinite medium under white-sky illumination.
- These are exactly the assumptions used to build the zero-variance walk, so it is the best case for what we can achieve.
- We can observe substantial variance reduction for low absorption, i.e. higher albedos (0.95 roughly corresponds to human skin).
- There is little, if any, variance reduction for lower albedos. This is due to the fact that the assumed shape of the solution does not match the true shape (see slide 27).
- However, when we take the computation time into account, we obtain a substantial net improvement in efficiency even in those cases because the resulting sub-surface walks are much shorter on average.
- Note that in the case that there is no absorption (albedo = 1), the classical subsurface walk is already zero variance: you just keep sampling (with whatever sample weight you had as you entered) until finally you exit: with that same weight no matter where you went, so there is nothing to improve upon in that case (unless, as mentioned above, you knew more about where light sources might enter the medium).


Use in rendering



- To use the method in rendering, we assume that the half-space is aligned with the surface normal at the point where the path enters under the surface.

Use in rendering – White sky illum.

25 samples per pixel




J. Křivánek - Zero variance walks for subsurface scattering

34

- The assumed half-space solution may substantially differ from the true solution on curved regions of the geometry.
- This may lead to occasional high sample weights, or fireflies (see top left image around the nose) because the estimated importance of each path was not accurately predicted.
- We mitigate this problem by using Multiple Importance Sampling (MIS) at each transition to combine classical and Dwivedi sampling.
- Note more fireflies near the nose and curved areas and their reduction as MIS mixes the classical and Dwivedi methods.
- Also note how the render time changes. While the use of the classical sampling is essential to avoid the fireflies for the curved regions, for flat regions like the cheek the pure Dwivedi walk works well and mixing in the classical sampling actually adds some noise. To get the best performance, the mixing weights could be driven by local curvature.

Use in rendering – IBL

equal-time comparison, 100 samples per pixel, 75% Dwivedi,



J. Křivánek - Zero variance walks for subsurface scattering


35

- Despite the fact that the importance function driving the sampling assumes uniform hemispherical illumination, the modified path sampling lowers variance even when the illumination is nonuniform.
- The images rendered with classical sampling use 100 samples/pixel while in our results we trace about 50% more samples/pixel in the same time.
- While the speedup is a profitable side-effect, most of the variance reduction is due to the sampling pdfs closely approximating the zero-variance sampling scheme.
- As on the previous slide, the subsurface medium has a single scattering albedo of 0.943 for all wavelengths and index-matched smooth boundaries.
- The individual rows show results for different environment maps.

WORK IN PROGRESS


Exponential radiance falloff

- ... does not hold for low scattering




- There are a few reasons for which the Dwivedi sampling scheme still yields a substantial amount of variance, despite being derived from a *zero-variance* scheme.
- First, the assumed solution inside of the medium is a poor approximation when absorption is high.

Directional distribution



- Second, the Dwivedi scheme assumes that the exponential importance solution extends all the way outside of the medium, when in fact, the importance is constant (equal to one) for any point outside of the medium (remember that 'escape' is our only source of importance and it does not matter to us where exactly we escape from the medium).

Radiance distribution



- We can see in these plots that there can be a significant discrepancy between the directional distributions assumed by Dwivedi and those obtained by a reference MC simulation, especially just below the surface. This is a consequence of both considering only the discrete asymptotic mode of the true solution and ignoring the boundary.

Improved sampling

- Take boundary into account
- Better radiance approximation
 - Matching 1st and 2nd moments of the true solution

- In our work on improving the Dwivedi scheme, we build on the paper by Hoogenboom [2008] who describes precisely how to construct a truly zero-variance walk for a half-space that considers the correct probability of escaping from the medium at any given step.
- To approach the zero variance ideal more closely, we employ a better approximation of the solution inside the half-space, based on matching the 1st and 2nd moments of the true solution.
- Another important aspect is that we explicitly take the boundary into account. This has a significant effect on the probability of escaping the medium and on the shape of the angular distribution.

Semi-infinite half-space test

albedo	0.4	0.8	0.95
improved scheme variance	 .0002 (100×)	 .0008 (116×)	 .001 (132×)
classical sampling variance	 0.020	 0.093	 0.132

- With this improved scheme, we are able to achieve a variance reduction of two orders of magnitude compared to classical sampling.
- So far, we have only tested the scheme in our synthetic half-space simulator. Application to rendering is an ongoing work.

Future work

- Boundary conditions (Fresnel, rough)
- Anisotropic scattering
- More rendering tests
- Other applications of zero-variance schemes

- Much work remains to be done.
- First, in our tests, we have assumed index-matched boundary. Changing the roughness or Fresnel ratio at the boundary does actually change the magnitude of the asymptotic term, but not its decay rate. Because the Dwivedi sampling discards the transient terms and renormalize, the magnitude of the asymptotic part goes away, and the Dwivedi scheme is then completely invariant to the boundary conditions.
- However, in our improved approach, changing the boundary condition - and therefore the magnitude of the solution - has a huge effect. This is because in the improved method we strictly separate the solution approximation inside of the medium (an *appropriately scaled* exponential), and outside (constant, unit source of importance). So one can say that Dwivedi's invariance to boundary conditions is only a lucky 'artifact' of some major simplifications in his scheme.
- Another open issue is anisotropic scattering. As the phase function changes away from isotropic, more and more discrete asymptotic terms appear. You can pick the largest one and still apply Dwivedi's scheme with success (we've tried HG with $g = 0.75$ and still get a nice variance reduction).
- Of course, we still need to apply our improved scheme in rendering.

- And finally, we believe there is much space for exploring other applications of the zero-variance theory in rendering problems.

Conclusion

- Solution for MC subsurface scattering
- Zero-variance MC schemes


- In conclusion, in this talk we have proposed a first practical solution for subsurface scattering based on MC simulation.
- Unlike for analytical BSSRDFs, the accuracy of the scheme is not affected by the departure from model assumptions.
- To achieve this goal, we have built upon the theory of zero-variance random walk. To the best of our knowledge, we are the first to bring this theory to the attention of the computer graphics research community.

Acknowledgments

- Funding from **Czech Science Foundation**

Zero-variance-based Monte Carlo subsurface scattering

Thank you!



Proof of zero variance (details)

$$\begin{aligned}
 \frac{1}{1 - P_{\text{term}}(\mathbf{r})} \left\langle \int L(\mathbf{r}') T(\mathbf{r}' \rightarrow \mathbf{r}) d\mathbf{r}' \right\rangle &= p(\mathbf{r}' | \mathbf{r}) = \frac{L(\mathbf{r}') T(\mathbf{r}' \rightarrow \mathbf{r})}{\int L(\mathbf{r}') T(\mathbf{r}' \rightarrow \mathbf{r}) d\mathbf{r}'} \\
 \frac{L(\mathbf{r})}{L(\mathbf{r}) - L^c(\mathbf{r})} \cdot \frac{\langle L(\mathbf{r}') \rangle T(\mathbf{r}' \rightarrow \mathbf{r})}{p(\mathbf{r}' | \mathbf{r})} &= \frac{L(\mathbf{r}') T(\mathbf{r}' \rightarrow \mathbf{r})}{L(\mathbf{r}) - L^c(\mathbf{r})} \\
 \frac{L(\mathbf{r})}{L(\mathbf{r}) - L^c(\mathbf{r})} \cdot \frac{\langle L(\mathbf{r}') \rangle T(\mathbf{r}' \rightarrow \mathbf{r})}{\left[\frac{L(\mathbf{r}') T(\mathbf{r}' \rightarrow \mathbf{r})}{L(\mathbf{r}) - L^c(\mathbf{r})} \right]} &= \\
 L(\mathbf{r}) \cdot \frac{\langle L(\mathbf{r}') \rangle}{L(\mathbf{r}')} &
 \end{aligned}$$