

Subdivision Next Event Estimation for Path-Traced Subsurface Scattering

David Koerner^{1,2,3}

Jan Novak¹

Peter Kutz²

Ralf Habel²


Wojciech Jarosz^{1,4}

¹Disney Research

²Walt Disney Animation Studios


³University of Stuttgart

⁴Dartmouth College



Hello. We present our work Subdivision Next Event Estimation for Path-Traced Subsurface scattering. This has been a joint effort with Wojciech Jarosz from Dartmouth College, Peter Kutz and Ralf Habel from the Disney Animation Studios and Jan Novak from Disney Research. My name is David Koerner and I will be presenting today.

Motivation




Photos licensed under Creative Commons BY-NC-ND

2

We deal with the problem of rendering translucent or subsurface scattering materials which are ubiquitous in everyday life and therefore are a common and important theme in rendering. At the same time, simulating this type of light transport is computationally very demanding, causing high cost in production. Improving the efficiency and performance for these computations is therefore an active research area of high demand in the industry.

Problem Statement




3

Next event estimation is a technique which improves performance for rendering surfaces and participating media, such as clouds. Although subsurface scattering is very similar to those effects, next event estimation can not be applied easily here. In our paper, we address this problem and propose a NEE technique for subsurface scattering materials which improves the performance of path-traced subsurface scattering.

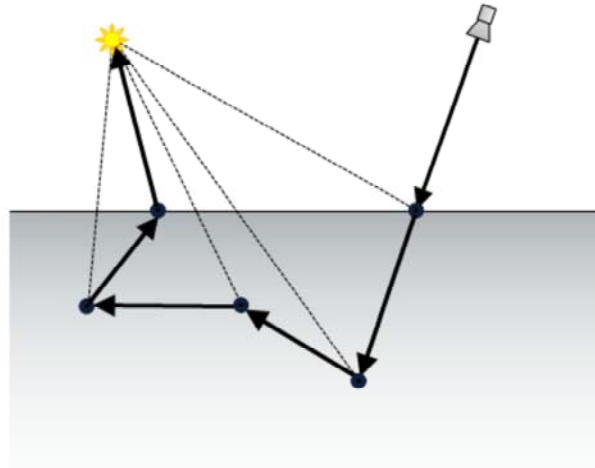
So what is next event estimation? Consider an infinite participating medium with a source and a sensor where unidirectional path tracing is used to connect both with a path sample.

Problem Statement



The idea of next-event estimation (NEE) or shadow rays is, to reuse the same path sample to sample light transport for different bounces. This is done by drawing explicit connections from every path vertex to the light source. This technique is very popular due to the performance improvements it brings.


Problem Statement



5

Now consider the presence of a boundary interface which is the case for all subsurface scattering objects.

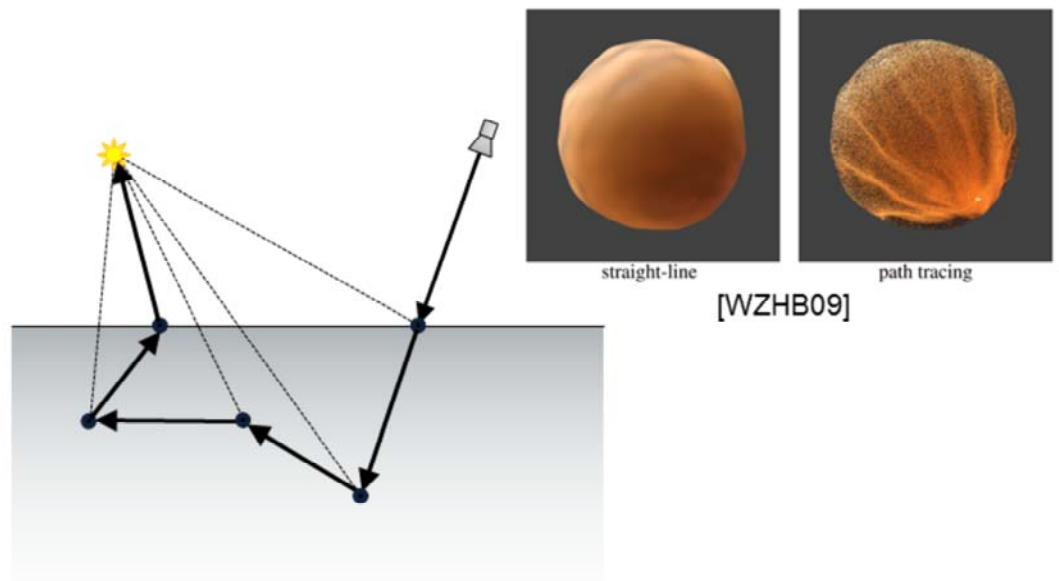
Problem Statement



6

Here light interacts with the material and refracts at the interface due to Snell's law of refraction. Therefore, direct connections are no longer valid which prevents straightforward application of NEE.

Problem Statement



7


An easy solution is to just ignore the surface interaction with the light. This gives the so called straight-line approximation. It has been shown by Walter and colleagues that this will miss important effects. Therefore, this shortcut is not an option for anyone who wants unbiased and accurate results.

Related Work

This problem of doing NEE through interfaces has been addressed in the past and still appears to be an active research topic.

Related Work

Single scattering in participating media with triangle mesh boundaries. [WZHB09]




9

First Walter et al. in 2009 presented „Single scattering in participating media with triangle mesh boundaries“. They proposed a brute force approach. First a camera ray sample is found <click> by random sampling.

Related Work

Single scattering in participating media with triangle mesh boundaries. [WZHB09]




10

Then for each triangle of the bounding mesh, a valid point on the triangle is found using a newton-raphson style iterative solver.

Related Work

Accurate computation of single scattering in participating media with refractive boundaries. [Hol15]




11

Last year, Holzschuch improved upon the work by Walter and his colleagues. He reversed the nesting of the inner loops in Walters approach. Instead of sampling the camera ray position first, they first go over all triangles and for each triangle, compute the interval of influence on the camera ray. Then the camera ray position is sampled within that interval using regular sampling. Given the camera ray position, newton-raphson is used again to find the final position on the triangle. This improves the convergence behaviour because regular sampling is now only used within piecewise smooth integrands per triangle along the camera ray.

Both methods only work for smooth dielectric interfaces.

Related Work

Manifold Next Event Estimation. [HDF15]




12

Manifold Next Event Estimation (MNEE) by Hanika et al. draws inspiration from perturbation strategies used in Markov Chain Monte-Carlo (MCMC) based rendering techniques such as MLT. They also follow the idea of finding a valid point on the surface. The key differences to the work by Walter and Holzschuch is that they represent the one bounce connection to the light in half vector space domain and use half vector space perturbation to find the point on the surface (HVSP).

With HVSP, the one-bounce connection is represented in half-vector space domain where paths samples are identified by their start and end point and a generalized half-vector per inbetween vertex. So the one-bounce connection is sampled by sampling such a half-vector <click>.

Related Work

Manifold Next Event Estimation. [HDF15]

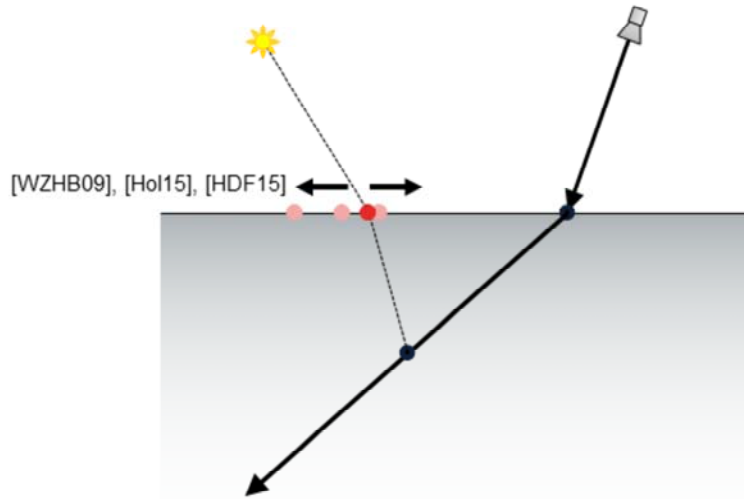


13

Then a newton style predictor-corrector scheme similar to manifold perturbation is used to find a valid point on the surface. The benefit is that this method does not need to rely on triangles as primitives for searching the points. And by using half vector space representation, their method supports different BSDFs and is not limited to smooth dielectrics.

However, in their paper, they actually dont apply MNEE to the problem of subsurface scattering although it theoretically could. The main reason for that is that the manifold walk becomes unreliable and unrobust with highly detailed geometry or displacement mapping.


Related Work



14

So common to all prior work is that idea of moving or finding a point on the surface which causes problems for highly detailed geometry. This is in contrast to our approach.


Idea



15

We try to investigate the idea of moving the last point within the medium to find a valid 2-bounce connection to the light. For this we want to use the degree of freedom given by the subsurface scattering volume.


Idea



16

Since our technique applies to every bounce within the volume, we now will just look at the general case where we just have a point in the medium which connects to the lightsource via a 2-bounce connection

Perturbation




17

Now consider standard uni-directional path construction starting from this point in the volume.

Perturbation

$$p(x_i x_j) = p(x_j | x_i) p(x_i)$$




18

During unidirectional path tracing, we then sample x_j which is found by directional and distance sampling at x_i .

Perturbation

$$p(x_i x_j x_k) = p(x_k | x_i x_j) p(x_j | x_i) p(x_i)$$




19

Then another point is sampled the same way which eventually hits the surface to form a one-bounce connection to the surface

Perturbation

$$p(x_i x_j x_k) = p(x_k | x_i x_j) p(x_j | x_i) p(x_i)$$




20

Now a lightray entering at x_k will refract into the volume and since x_j will not line up perfectly, the path will not receive any light. The x_i, x_j, x_k connection will be invalid for the given incoming light direction.

Now following the perturbation approach of previous work, the idea would be to just move x_j into a correct position.

Perturbation


$$p(x_i x_j x_k) = p(x_k | x_i x_j) p(x_j | x_i) p(x_i)$$



21

Ideally we would just like to move or resample x_j such that it lies on the refracted ray and therefore constitutes a valid path.

Perturbation




$$p(x_i x_j x_k) = p(x_k | x_i x_j) p(x_j | x_i) p(x_i)$$

$$p(x_i x'_j x_k) = p(x'_j | x_i x_k) p(x_i x_k)$$

22

This step needs to be reflected in the pdf. The probability of x'_j is conditioned on vertices x_i and x_k .

Perturbation



$$p(x_i x_j x_k) = p(x_k | x_i x_j) p(x_j | x_i) p(x_i)$$

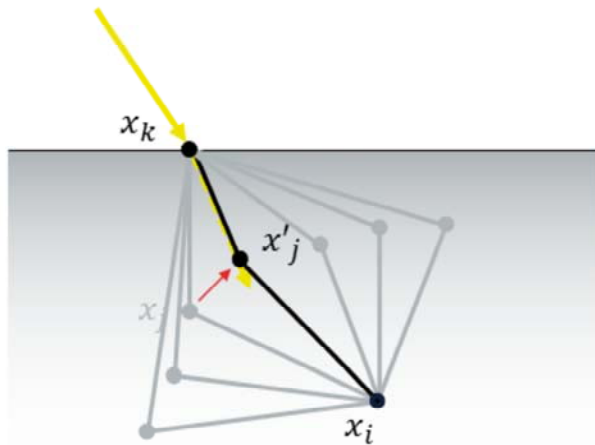
$$p(x_i x'_j x_k) = p(x'_j | x_i x_k) p(x_i x_k)$$

$$p(x_i x_k) = \int_{x_j \in \mathcal{V}} p(x_i x_j x_k) dx_j$$

23

This probability needs to be found by marginalization...

Perturbation



$$p(x_i x'_j x_k) = p(x'_j | x_i x_k) p(x_i x_k)$$

$$p(x_i x_k) = \int_{x_j \in \mathcal{V}} p(x_i x_j x_k) dx_j$$

Difficult!

24

...which is impractical to compute as we need to integrate over the space of all two segment paths that connect x_i and x_k .

MNEE faces the same problem when perturbing vertices on the surface. Their solution is to have a bijective mapping from x_j to x'_j . For a given x_j , there is only a single valid x'_j found by their technique (one-to-one mapping). In that case the integrand is a delta distribution and the integral collapses.

We don't have such a bijective mapping. Our approach is to circumvent this problem by avoiding marginalization altogether.

Subdivision

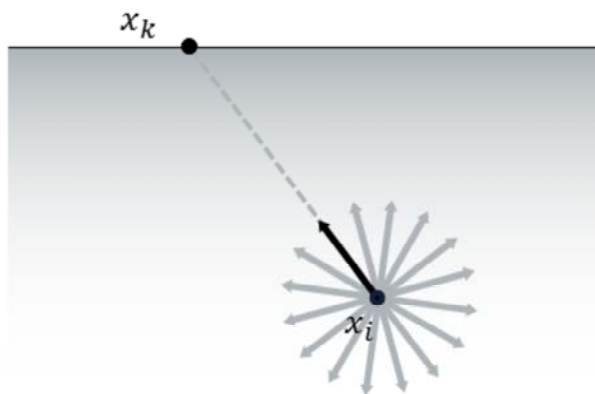


25

For this we use subdivision.

The main idea behind our method is to simply change the order in which path vertices are sampled. Incremental path construction is often understood as extending the path on one of its ends with a new vertex. However, it is perfectly valid to sample the path vertices in any arbitrary order, which means that the path can be extended at any of its intermediate edges. By adding vertices between two existing vertices of the path, we basically subdivide the edge they share. This is where the name of our method comes from.


Subdivision



26

So after arriving at x_i , we next sample x_k on the surface. Since we need to have a surface point, we need to have a method for explicitly sampling surface vertices. We use the very simple approach of randomly sampling a direction at x_i and taking the first intersection. Since we are within a bounded volume, it is guaranteed that there is one.

Subdivision




27

THEN we sample x_j . Since we now know the surface point we can use information about the bsdF and incoming light when sampling x_j .

Subdivision

$$p(x_i x_k) = p(x_k | x_i) p(x_i)$$




28

So the pdf for segment x_i to x_k is the pdf for x_k conditioned on pdf for x_i .

Subdivision

$$p(x_i x_k) = p(x_k | x_i) p(x_i)$$


$$p(x_i x_j x_k) = p(x_j | x_i x_k) p(x_i x_k)$$



29

When sampling x_j , its pdf is conditioned on the pdf for x_i and x_k . It simply is a conditional probability. Instead of having to do marginalization we basically add another dimension to the pdf measure and the existing dimensions can be left untouched.

Subdivision




$$p(x_i x_k) = p(x_k | x_i) p(x_i)$$
$$p(x_i x_j x_k) = \underbrace{p(x_j | x_i x_k)}_{?} p(x_i x_k)$$

30

The remaining question is: How do we sample x_j ?

We find a point in the volume by sampling a direction and distance. Since we can do that either by starting from x_k or x_i , we have two ways to sample x_j , where each takes different terms in the rendering equation into account.


Subdivision – Top To Bottom



31

The first we refer to as top-to-bottom subdivision. We sample a direction to the light source at x_k ...


Subdivision – Top To Bottom



32

...followed by sampling the BTDF part of the BSDF at x_k . This gives a direction into the volume which has been sampled according to incoming light and BSDF at x_k .


Subdivision – Top To Bottom



33

Then we sample a distance which will give us x_j . From there we make a direct connection to x_i to close the path. We use the one-sample MIS estimator to combine transmittance and equiangular sampling in order to account for transmittance and geometry term during distance sampling.


Subdivision – Bottom To Top



34

The second strategy we call Bottom To Top subdivision.

Subdivision – Bottom To Top




Joint importance sampling [GKHJ13]

35

Here we use (single bounce) joint importance sampling to sample direction and distance according to phase function at x_i and geometry term. This gives x_j .

Then we make a direct connection from x_j to x_k and sample the BTDF part of the BSDF at x_k to get an outgoing direction according to the BSDF.



Subdivision



36

Our two different subdivision strategies give us two different sub-paths. We use the MIS multi-sample estimator to combine both contributions and therefore account for different terms in the rendering equation.


Subdivision



37

It is interesting to realize the similarity to surface rendering where at every bounce, a bsdf and light sample is drawn which both are combined using MIS. Our technique can be understood as extending this very same concept to the first bounce within the volume. In addition to a direction at x_k , our light or bsdf samples consist of an additional vertex or 2-segment subpath within the volume.

Subdivision



38

Another interpretation of our method can be seen by realizing that our method resembles one-bounce bi-directional pathtracing within a single edge of our parent path. The difference to standard BDPT is that start end end points are not source and sensor and that we only sample a single bounce. In our discussions, we sometimes would refer to this as local bi-directional path tracing.

Simply changing the order in which path vertices are constructed leads to these different perspectives and interpretations.

Subdivision Next Event Estimation (SNEE)

At each path vertex x_i :


1. Find exit point x_k on the surface
2. Subdivide and place x_j
 - Using top-to-bottom subdivision
 - Using bottom-to-top subdivision
3. Combine using multi-sample MIS

39

To summarize, our NEE method has the following parts:

1. Find an exit point x_k on the surface
2. Subdivide the edge which connects that exit point with the current path vertex.
There are 2 strategies which give two different samples.
3. Combine those using MIS

Assumptions



40

Our technique has some assumptions which we need to mention:

1. x_i has to be in the volume.
2. x_k must be on the surface and directly visible from x_i .
3. Directional degree of freedom at x_i and x_k is required to be able to freely place x_j within the volume.

Unhandled Transport

41

Since our assumptions are not always met, our method does not sample all light transport.

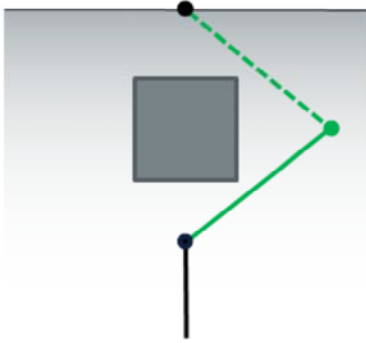
Unhandled Transport



42

The transport which is not handled is the transport where the last two segments within the medium have an invalid parent edge. Consider this case of an occluder within the medium.


Unhandled Transport



43

Further consider this path with the following one bounce connection to the surface.


Unhandled Transport



44

In our technique the only way this transport would be sampled is by subdividing the edge shown in grey.


Unhandled Transport



45

However, this edge will never exist due to the way we currently sample the surface points. Therefore it will never be subdivided and the shown transport would never be sampled. A similar problem arises for parent edges on the flat surface. Those will never be sampled.


Unhandled Transport



46

A similar problem arises for parent edges on the flat surface. Those will never be sampled.

Unhandled Transport



47

When both endpoints are on a surface and we have a smooth dielectric bsdf, there is no degree of freedom for choosing a direction at each point and hence our technique is not able to sample an intermediate vertex for the one-bounce connection in this case.

In our current implementation, we exclude all transport, where one of the last two vertices within the medium is on a surface. This obviously includes single scattering.

Integration

To account for unhandled transport, we integrate SNEE with a standard unidirectional path tracer. The way we do it is by interpreting the last two-segments of our standard path tracing sample as a one-bounce NEE technique and use multi-sample MIS to combine it with SNEE.


Integration



49

Consider this case where we are within the volume during unidirectional path tracing.


Integration



50

At every bounce we will do SNEE which gives us two one-bounce samples from the two different sampling strategies we have.


Integration



51

Then we continue incrementally building the path and doing SNEE at every bounce (not showing the second SNEE here).


Integration



52

Eventually we will hit the surface, which finishes our random walk through the medium.

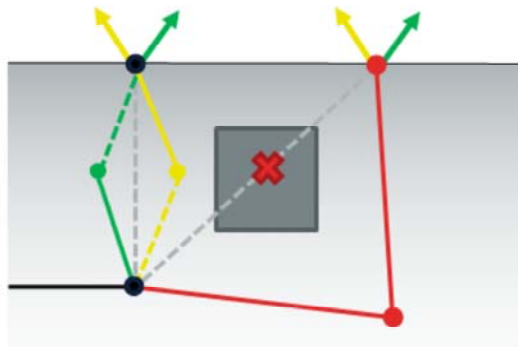
Integration



53

The last two segments of our random walk can be understood as a special next event estimation technique, which uses standard incremental path construction to find a one-bounce connection to the surface.

Integration




55

Fail cases are naturally handled by setting the pdf of SNEE to zero for samples which it can not handle.

Results

We now discuss the main result of our paper.

Results



57


In order to study the behaviour of our technique, we exposed it to different configurations of illumination and bsdf. Here we have a directional light source and the first column uses a rough dielectric interface with a high roughness value of 0.2. The left half in each image has been rendered using standard path tracing with MIS between bsdf and light sampling. The right halves use SNEE. These are equal time comparisons.

As we can see in the very left column. Our method does not provide any benefit for the case of a very rough surface with a weak directional preference. This is because when sampling the light, the probability of hitting the lobe of the bsdf is high.

When we increase the directionality of the surface [click](#) we start to see improvement. Here we use a small roughness value of 0.04 which gives the interface a strong directional preference. We can see that our method starts outperforming the standard method.

In case of a purely directional interface such as a smooth dielectric and a delta light illumination [click](#), standard path tracing is not able to produce a solution at all, while our method comes up with a result.

Results





58

Now we start increasing the directionality of the light source. Here we use a solid angle light with a solid angle of about 100x the solid angle subtended by the sun.

For the high roughness bsdf we only see a very minimal improvement and as we decrease the roughness value, our method outperforms standard sampling again. As the bsdf approaches a delta distribution the standard solution is increasingly poor.

Results



59

Now when we use a low frequency environment light illumination, we see that our method does not improve for BxDF's with any directional preference. This is because just using bsdf samples is enough to sample the low frequency lightsource.

Future Work

- Evaluation
- Study surface point sampling and russian roulette to improve the method

60

...which brings us directly to the various things we need to do onwards.

The work we presented today is very much work in progress. All reviewers asked for more extensive evaluation which we very much agree with and the lack of comparisons is the reason why we submitted to the experimental ideas and implementations track of egsr. We are planning to do comparisons with MNEE and the work from Holzschuch for the smooth dielectric. <click>

Currently our method only has a benefit in cases of strong directionality of light or bsdf and degenerates for uniform illumination and surface scattering. It needs to be made more general in order to be of practical use.

Related to the previous point is the fact that we want to investigate more surface point sampling techniques and the use of russian roulette to improve the performance of our method in the more general cases. <click>

Conclusion

- Being creative in path construction is fruitful
- Mutation and perturbation strategies in MCMC inspire new methods in other rendering techniques

61

While SNEE is probably not practical in its current form, our work shows that being creative in the way path samples are constructed can be beneficial.

Mutation and perturbation strategies in MCMC methods can inspire new avenues in other rendering techniques.

Q&A

62

Thank you for your attention. We are happy to take questions.

Bibliography

[GKH*13] Georgiev I., Krivanek J., Hachisuka T., Nowrouzezahrai D., Jarosz W.: Joint importance sampling of low-order volumetric scattering. *ACM TOG (Proc. of Siggraph Asia)* 32, 6 (2013), 164:1-164:14.

[HDF15] Hanika J., Droske M., Fascione L.: Manifold next event estimation. *CGF (Proc. of Eurographics Symposium on Rendering)* 34, 4 (2015), 87-97.

[Hol15] Holzschuch N.: Accurate computation of single scattering in participating media with refractive boundaries. *CGF (Proc. of Eurographics Symposium on Rendering)* 34, 6 (2015), 48-59.

[WZHB09] Walter B., Zhao S., Holzschuch N., Bala K.: Single scattering in participating media with triangle mesh boundaries. *ACM Transaction on Graphics (Proc. of Siggraph 2009)* 28, 3 (2009), 92:1-92:8.