Eurographics 2008

Practical Product Importance Sampling for Direct Illumination

Petrik Clarberg Tomas Akenine-Möller

Lund University
Sweden

This work was presented by Petrik Clarberg at Eurographics 2008 in Crete, on April 18, 2008.
Comments have been added afterwards. Please send any questions to: petrik@cs.lth.se.

Motivation

  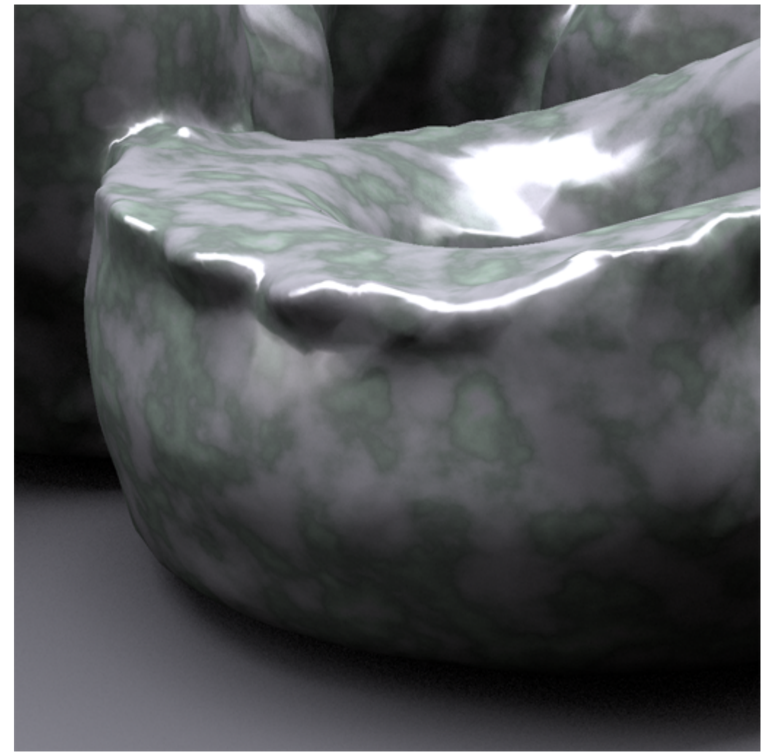

Monte Carlo ray tracing of direct illumination

- Environment map (and/or distant area) lighting
- Realistic materials

General, easy to use method

We focus on stochastic ray tracing of direct illumination, using importance sampling of the product of distant lighting and surface reflectance. Our goal is a method that is simple and easy to implement, and works for arbitrary BRDFs.

The Problem

Direct Illumination

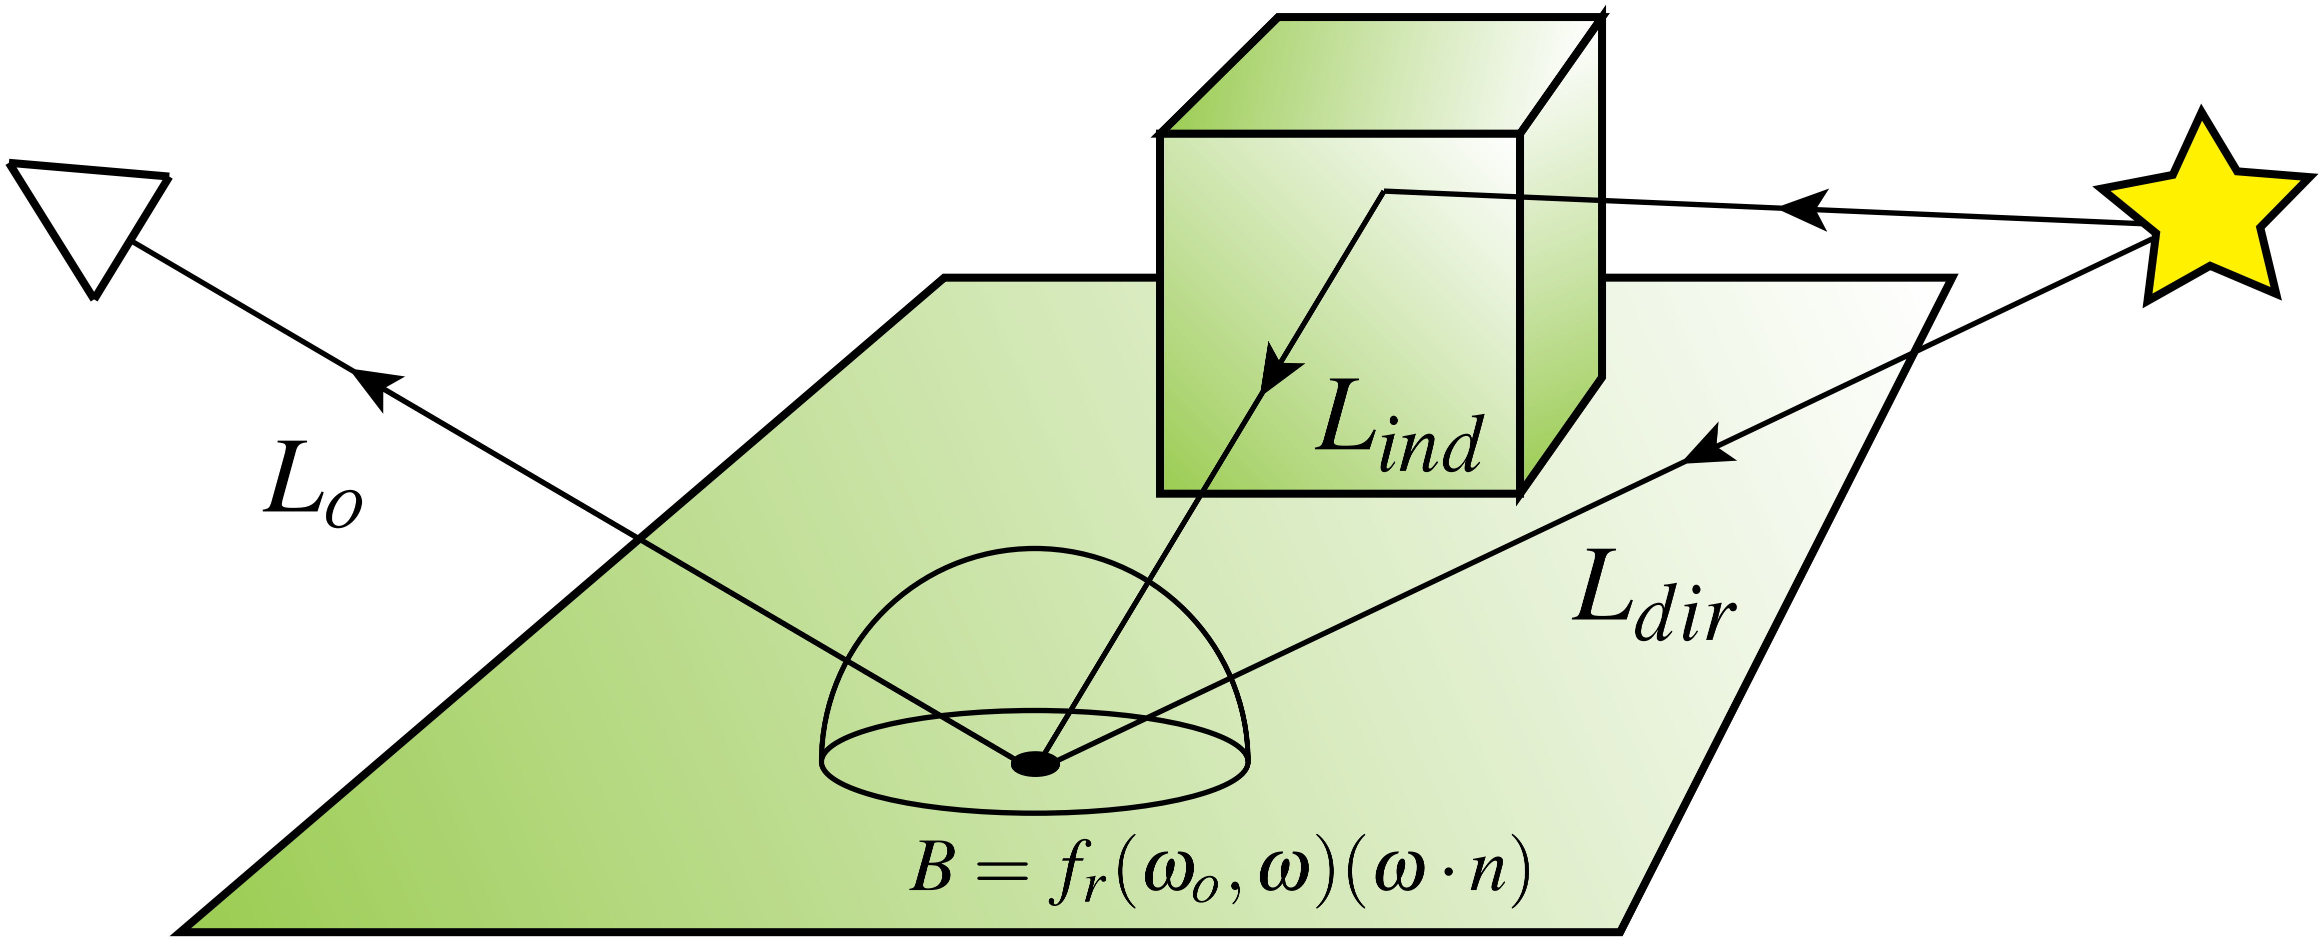

$$L_o = \int L\, B\, d\omega = \int L_{dir}\, B\, V\, d\omega + \int L_{ind}\, B\, d\omega$$

Practical Product Importance Sampling for Direct Illumination, Petrik Clarberg and Tomas Akenine-Möller, EG 2008

The rendering equation [Kajiya86] gives the outgoing radiance, L_o, as an integral over the incident lighting L, and reflectance B (BRDF x cosine-term). This can be split into indirect illumination (one or more bounces) and direct illumination (zero bounces). The integral for the direct illumination involves a binary visibility term V.

Direct Illumination

$$L_o = \int L_{dir}\, B\, V\, d\omega$$

Point lights: $L_o = \sum_i L_i\, B\, V$

Area lights

- N random samples per light

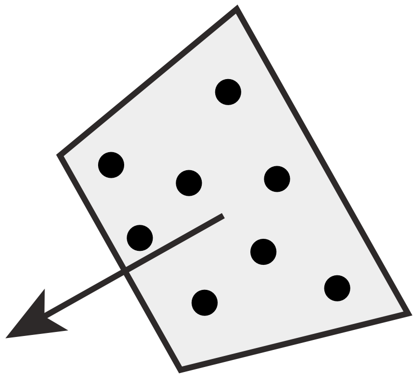

Environment map lighting

- Divide into pre-integrated strata: e.g., [Agarwal et al. 2003, Ben-Artzi et al. 2006]
- Lightcuts [Walter et al. 2005, 2006]
- Monte Carlo simulation

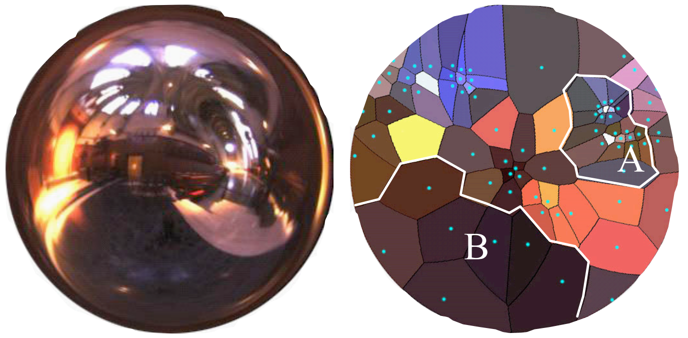

Image courtesy of Ben-Artzi et al.

Depending on what type of light sources are used, there are different ways of computing the direct illumination. We focus on Monte Carlo methods, which rely on stochastic point sampling of the integral.

Monte Carlo Integration

Random sampling:

$$I = \int f(x)\,dx \approx \frac{1}{N}\sum_i f(x_i)$$

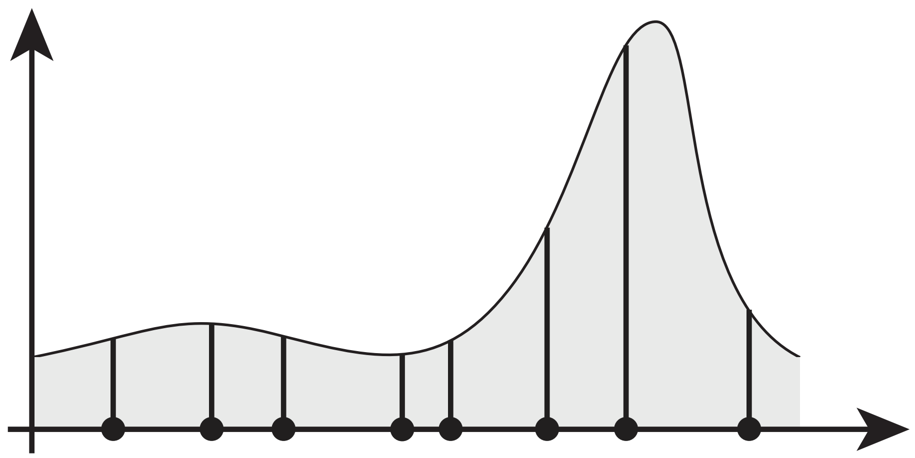

Importance sampling:

$$I \approx \frac{1}{N}\sum_i \frac{f(x_i)}{p(x_i)}$$

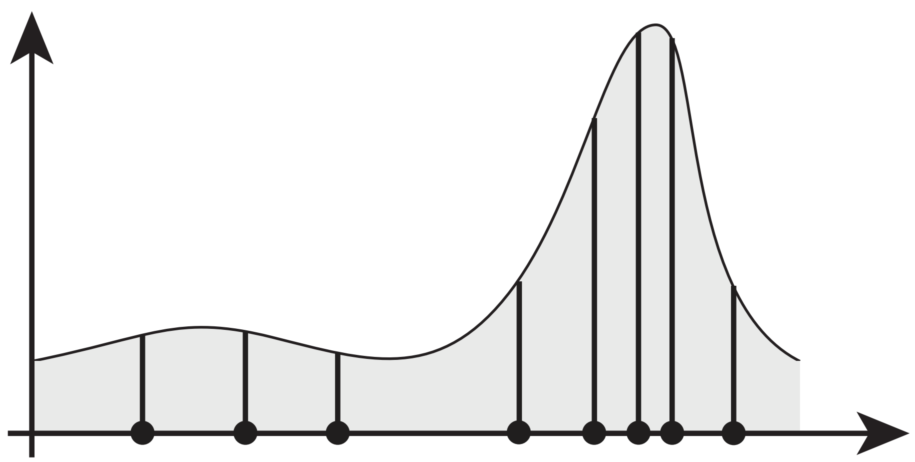

Zero variance if $p(x_i) \propto f(x_i)$

Practical Product Importance Sampling for Direct Illumination, Petrik Clarberg and Tomas Akenine-Möller, EG 2008

Random sampling gives noisy results as the samples have a low probability of finding the peaks in the function. With importance sampling, more samples are placed in regions of high contribution, thereby reducing the variance. The optimal case is sampling with a PDF proportional to the function f(x), but this is impossible to achieve in practice.

Product Sampling

Problem with sampling only *L*, *B*, or *V* ?

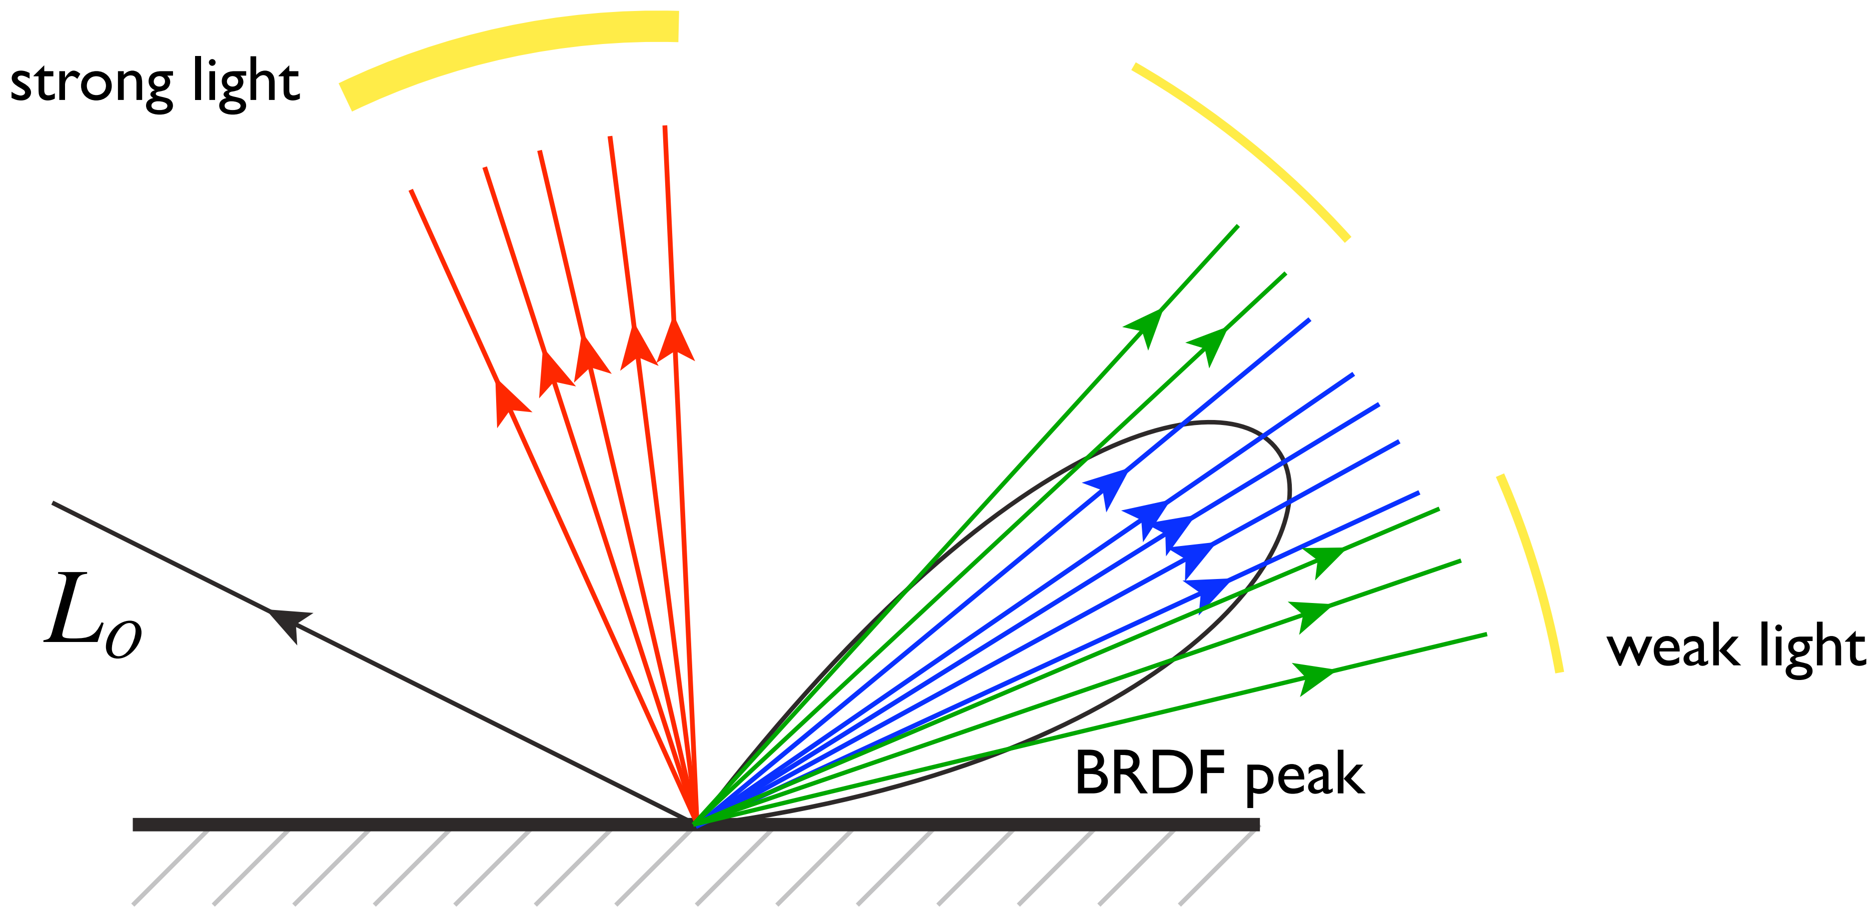

Product sampling: $L \cdot B$

- Find samples with <u>both</u> strong light and BRDF peak

Practical Product Importance Sampling for Direct Illumination, Petrik Clarberg and Tomas Akenine-Möller, EG 2008

With BRDF sampling (blue), the samples are placed in the direction of the BRDF peak, which in this case means they miss all important lighting. Similarly, sampling according to the lighting only (red) may give bad results since we BRDF is not taken into account. Product sampling (green) solves the problem.

Problem

$p(\omega) =$? x ?

Lighting BRDF

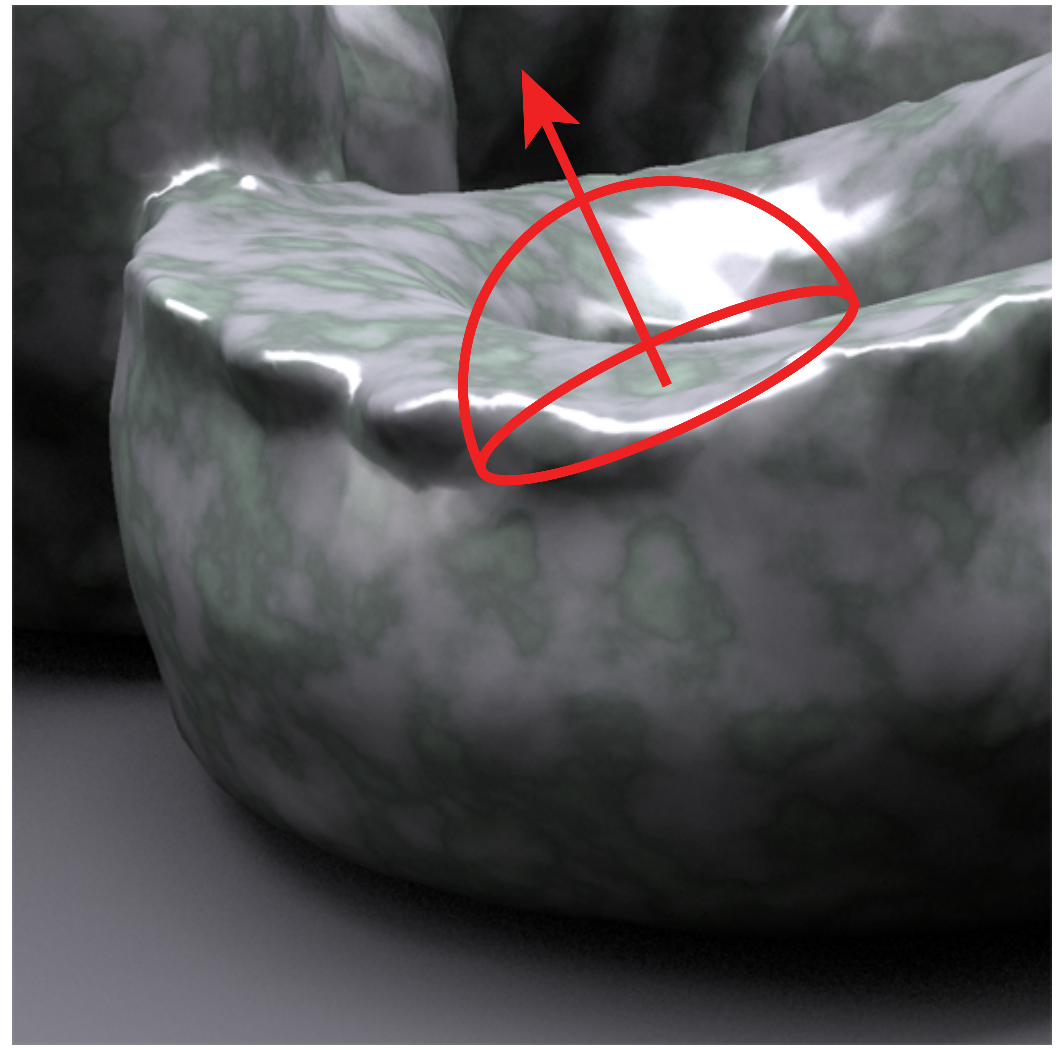

BRDF typically changes per pixel

Lighting (local frame) also changes per pixel

How to compute and sample the product?

Practical Product Importance Sampling for Direct Illumination, Petrik Clarberg and Tomas Akenine-Möller, EG 2008

The probability density function (pdf) changes per-pixel, so cannot easily be precomputed. The BRDF is a 4D function in the local frame (and typically spatially varying), and distant lighting is a 2D function in world space (5D in local frame). This makes efficient computation and sampling of the product a difficult problem.

Wavelet Imp. Sampling (WIS)

[Clarberg et al. 2005]

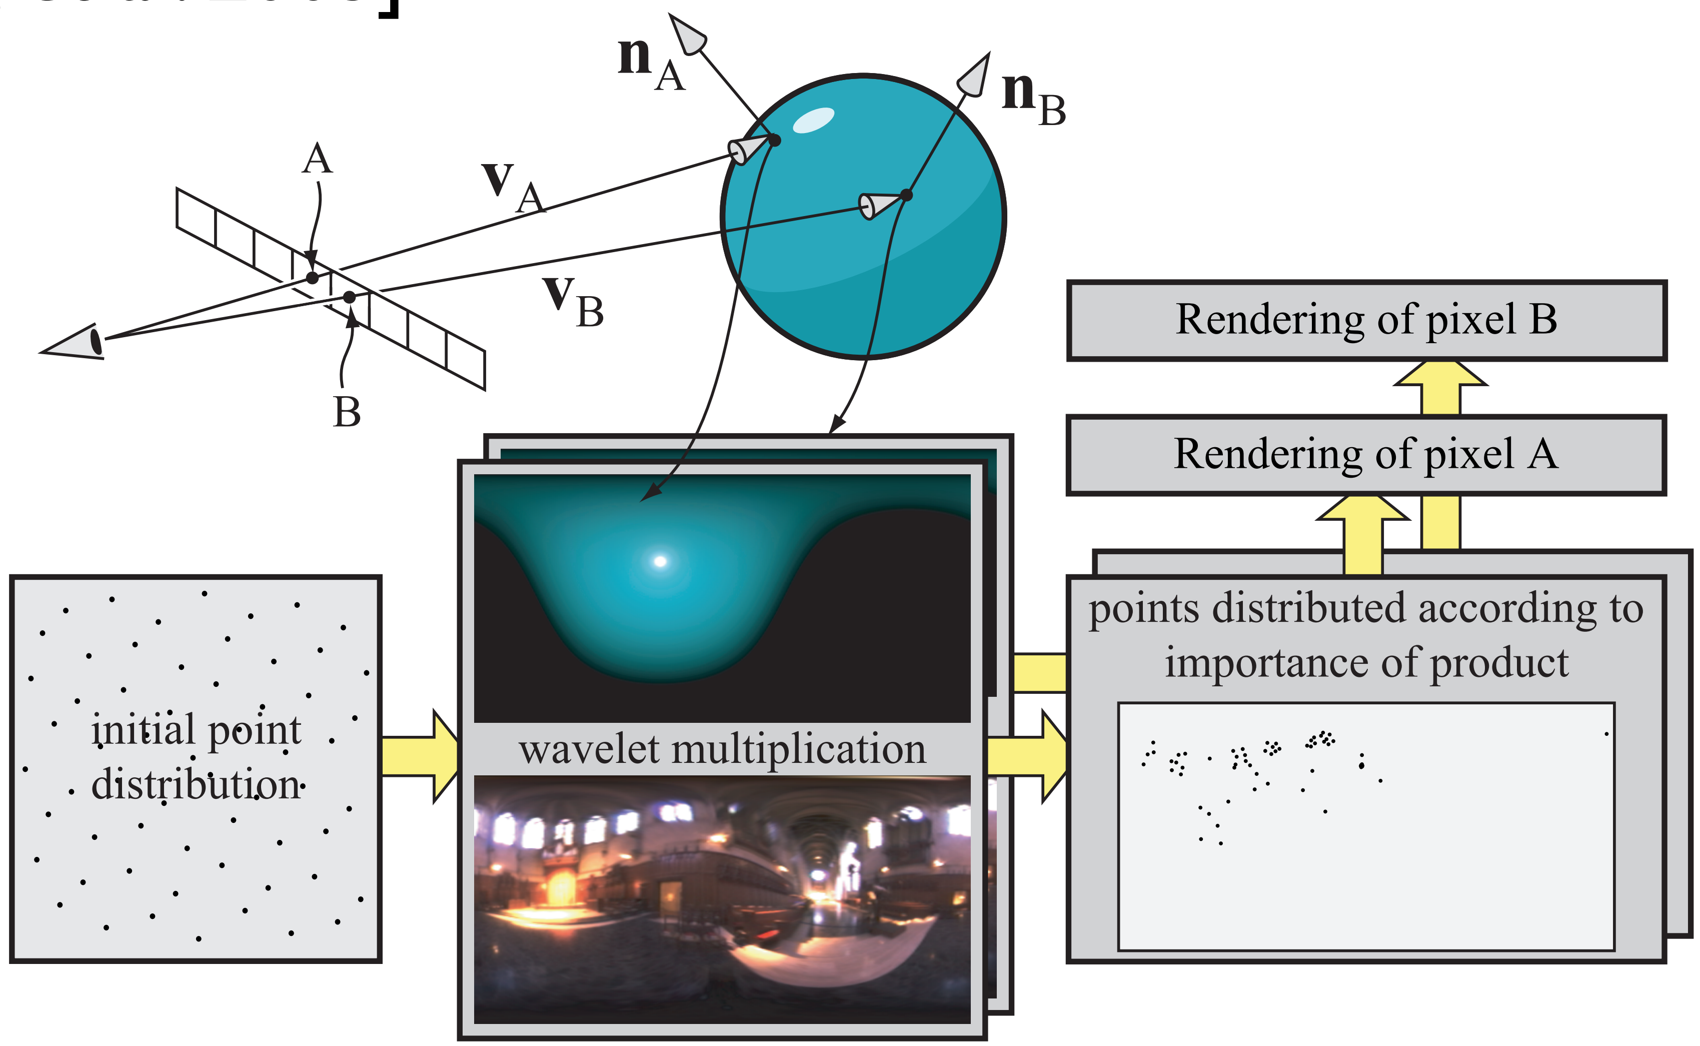

$L \cdot B$

Limitations

× Precomputation, low-res lighting (4D pre-rotation)

× Adding dimensions ⟶ Exploding memory & time

WIS solves the problem by precomputing and storing L and B as Haar wavelets. These are then multiplied (see: wavelet triple products [Ng et al. 2004]) and sampled using sample warping. However, the large amounts of data limits the method to low-res lighting and simple BRDFs (i.e., it is hard to use procedural shaders of spatially-varying BRDFs).

Two Stage Imp. Sampling [Cline et al. 2006]

Summed area table (SAT) of Env. map

- Quick to find integral over arbitrary region

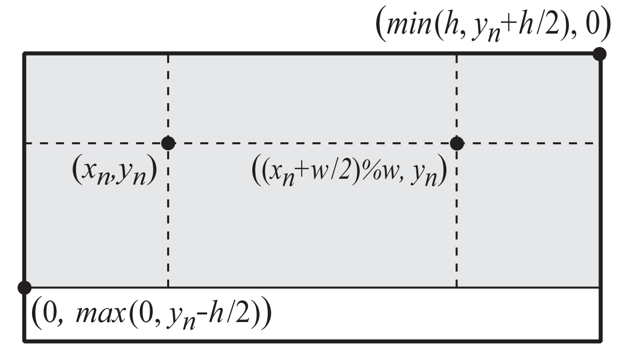

Piecewise linear BRDF approx.

- Initial splits based on BRDF peaks:
- Heuristics control further splits:
- × Relies on heuristics
- × Relatively inaccurate

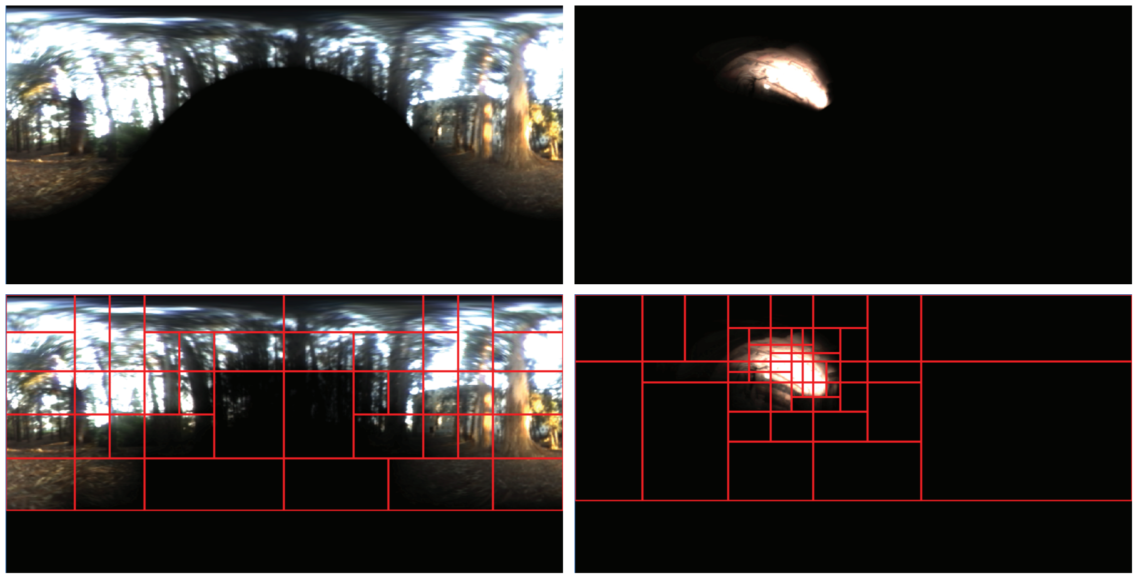

Practical Product Importance Sampling for Direct Illumination, Petrik Clarberg and Tomas Akenine-Möller, EG 2008

Cline et al. solved many of the problems by using a SAT for the environment map in world space, and building a BRDF approximation on-the-fly per-pixel. The main drawback is that the construction of the BRDF approximation relies on heuristics and requires a rather large number of splits.

Our Solution

Lighting

Do we need full-res environment map?

- Subsampling / approx. would simplify the problem

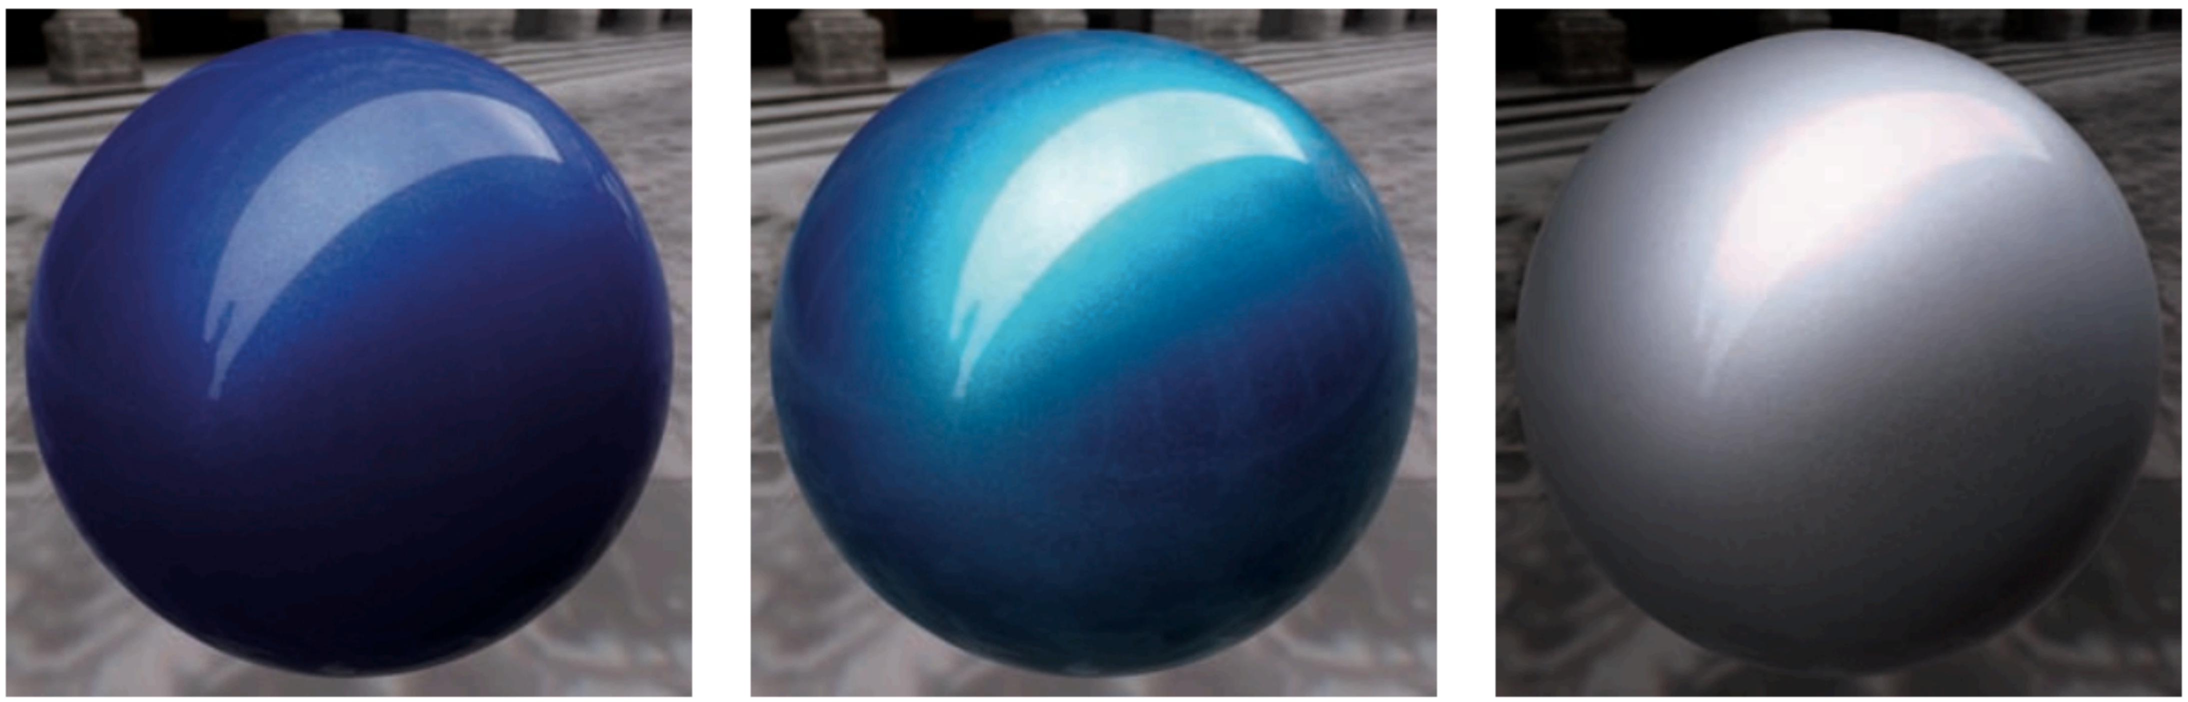

Image courtesy of Rump et al. 2008

Practical Product Importance Sampling for Direct Illumination, Petrik Clarberg and Tomas Akenine-Möller, EG 2008

These renderings of realistic car paint represent a case where the BRDFs have both a wide diffuse lobe, and a very sharp specular peak. For these types of materials, any kind of subsampling of approximation of the lighting yields significantly increased noise, as the BRDF peak can point towards any region of the environment map.

Lighting

Do we need full-res environment map? **Yes**

4K HDR map = 192 MB (uncompressed)

Pre-rotation (4K x 4K directions): 3072 TB

Conclusion

- Single map in world space
- Mip-map hierarchy (+33% memory)
- Which format?

So, yes, it is in many cases desirable to use full-res environment maps. The industry often uses maps of 4Kx4K pixels, so pre-rotating (as in WIS) is clearly impractical (>3000TB memory!). Hence, we must store a single map in world space.

Equal-Area Mapping

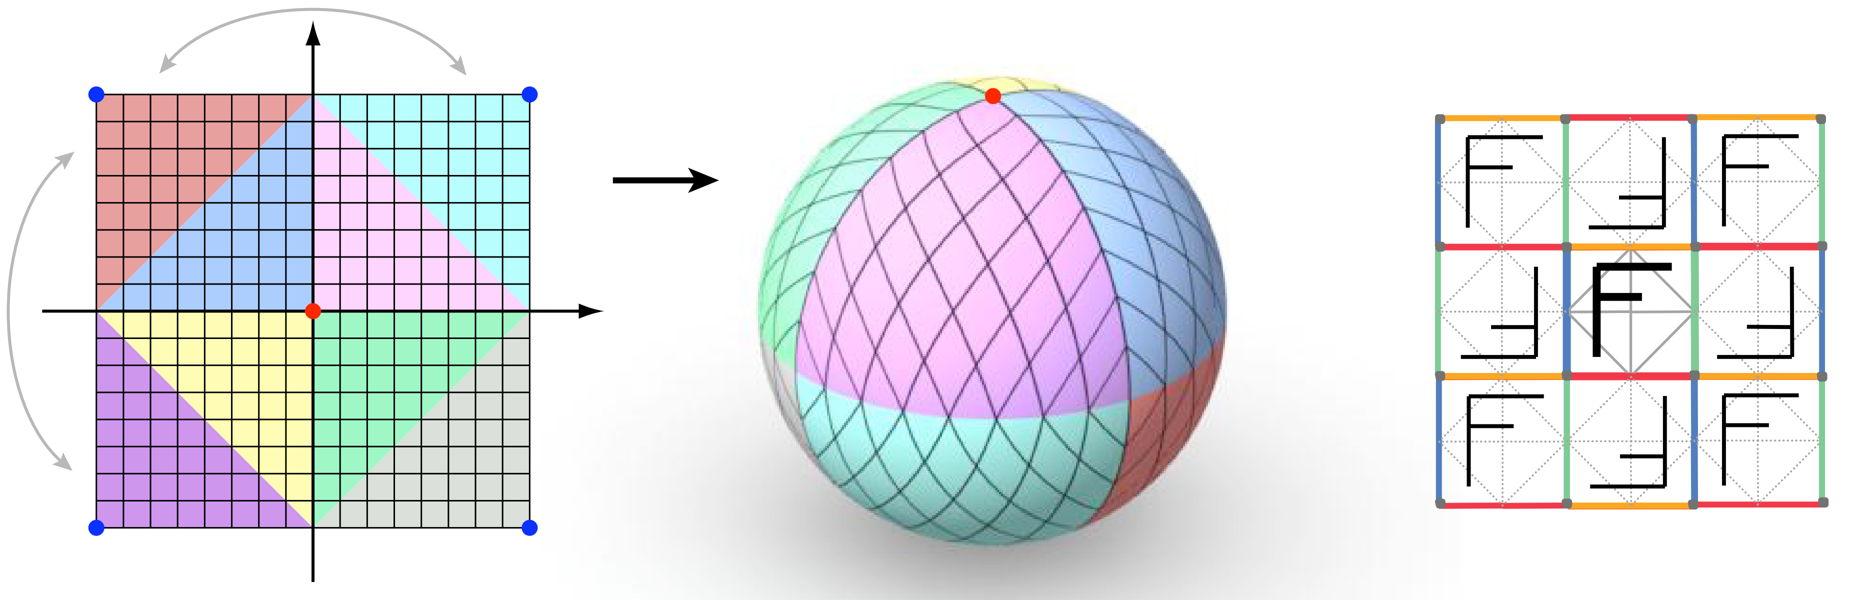

(a) (b) (c)

Properties

- Equal-area, Low distortion, Boundary symmetry

Speed

Scalar	120 clocks/transform
SSE	20 clocks/transform

Practical Product Importance Sampling for Direct Illumination, Petrik Clarberg and Tomas Akenine-Möller, EG 2008

We have constructing an equal-area mapping of the sphere by combining the concentric map [Shirley and Chiu 97] with the octahedral map [Praun and Hoppe 2003]. We have written a fast SSE-optimized version of the mapping, which will be described in a separate paper or technical report. (For the EG'08 paper, we used the slower scalar version).

BRDF Approximation

$$p(\omega) =$$ 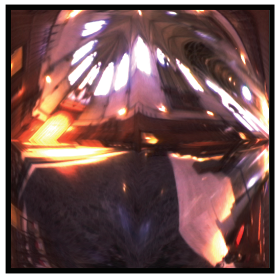 x ?

Lighting   BRDF

Need BRDF slice in world space

- 4D BRDF → 7D function
- 6D BTF → 9D function

Precomputation impractical

- Create approximation on-the-fly

Now when the lighting has been defined, how do we find the BRDF slice? Pre-rotation is impractical (7D or even 9D function), so we must create an approximation on-the-fly per pixel.

BRDF Approximation

Draw point samples

- Evaluate BRDF x cosine for each sample
- Use BRDF importance sampling to find important features

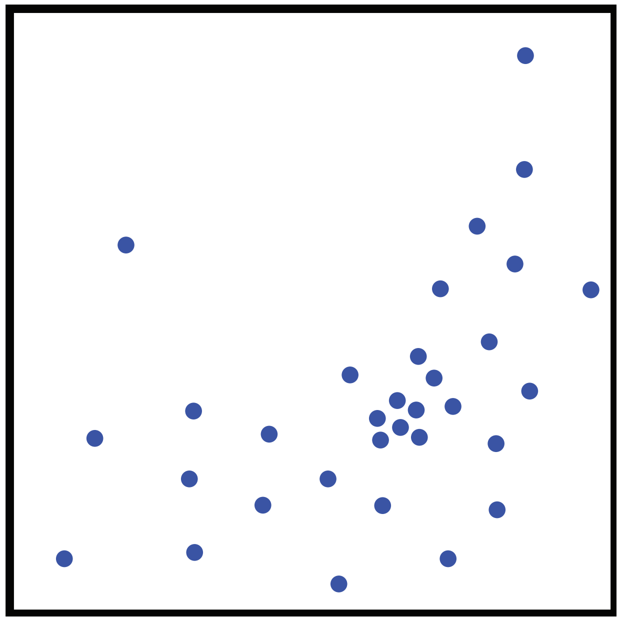

Hierarchically subdivide samples into quadtree

- One sample per node
- Piecewise constant approx.
- High resolution around peak(s)

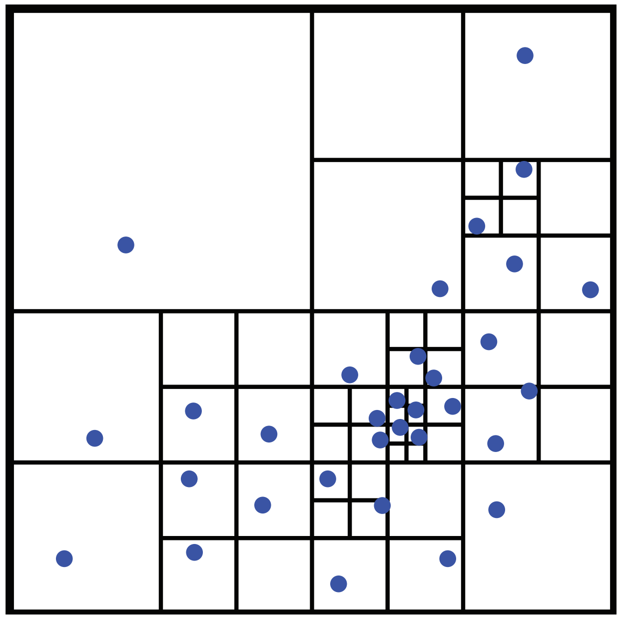

Practical Product Importance Sampling for Direct Illumination, Petrik Clarberg and Tomas Akenine-Möller, EG 2008

The core idea is to build an approximation out of a small set of point samples of the BRDF. We use BRDF importance sampling to direct the samples to important regions. Based on the point samples, a piecewise constant approximation is built by subdividing the sample set in a quad-tree structure. Pseudo-code for this can be found in the paper.

Quadtree Product

See paper for Fast Wavelet Product

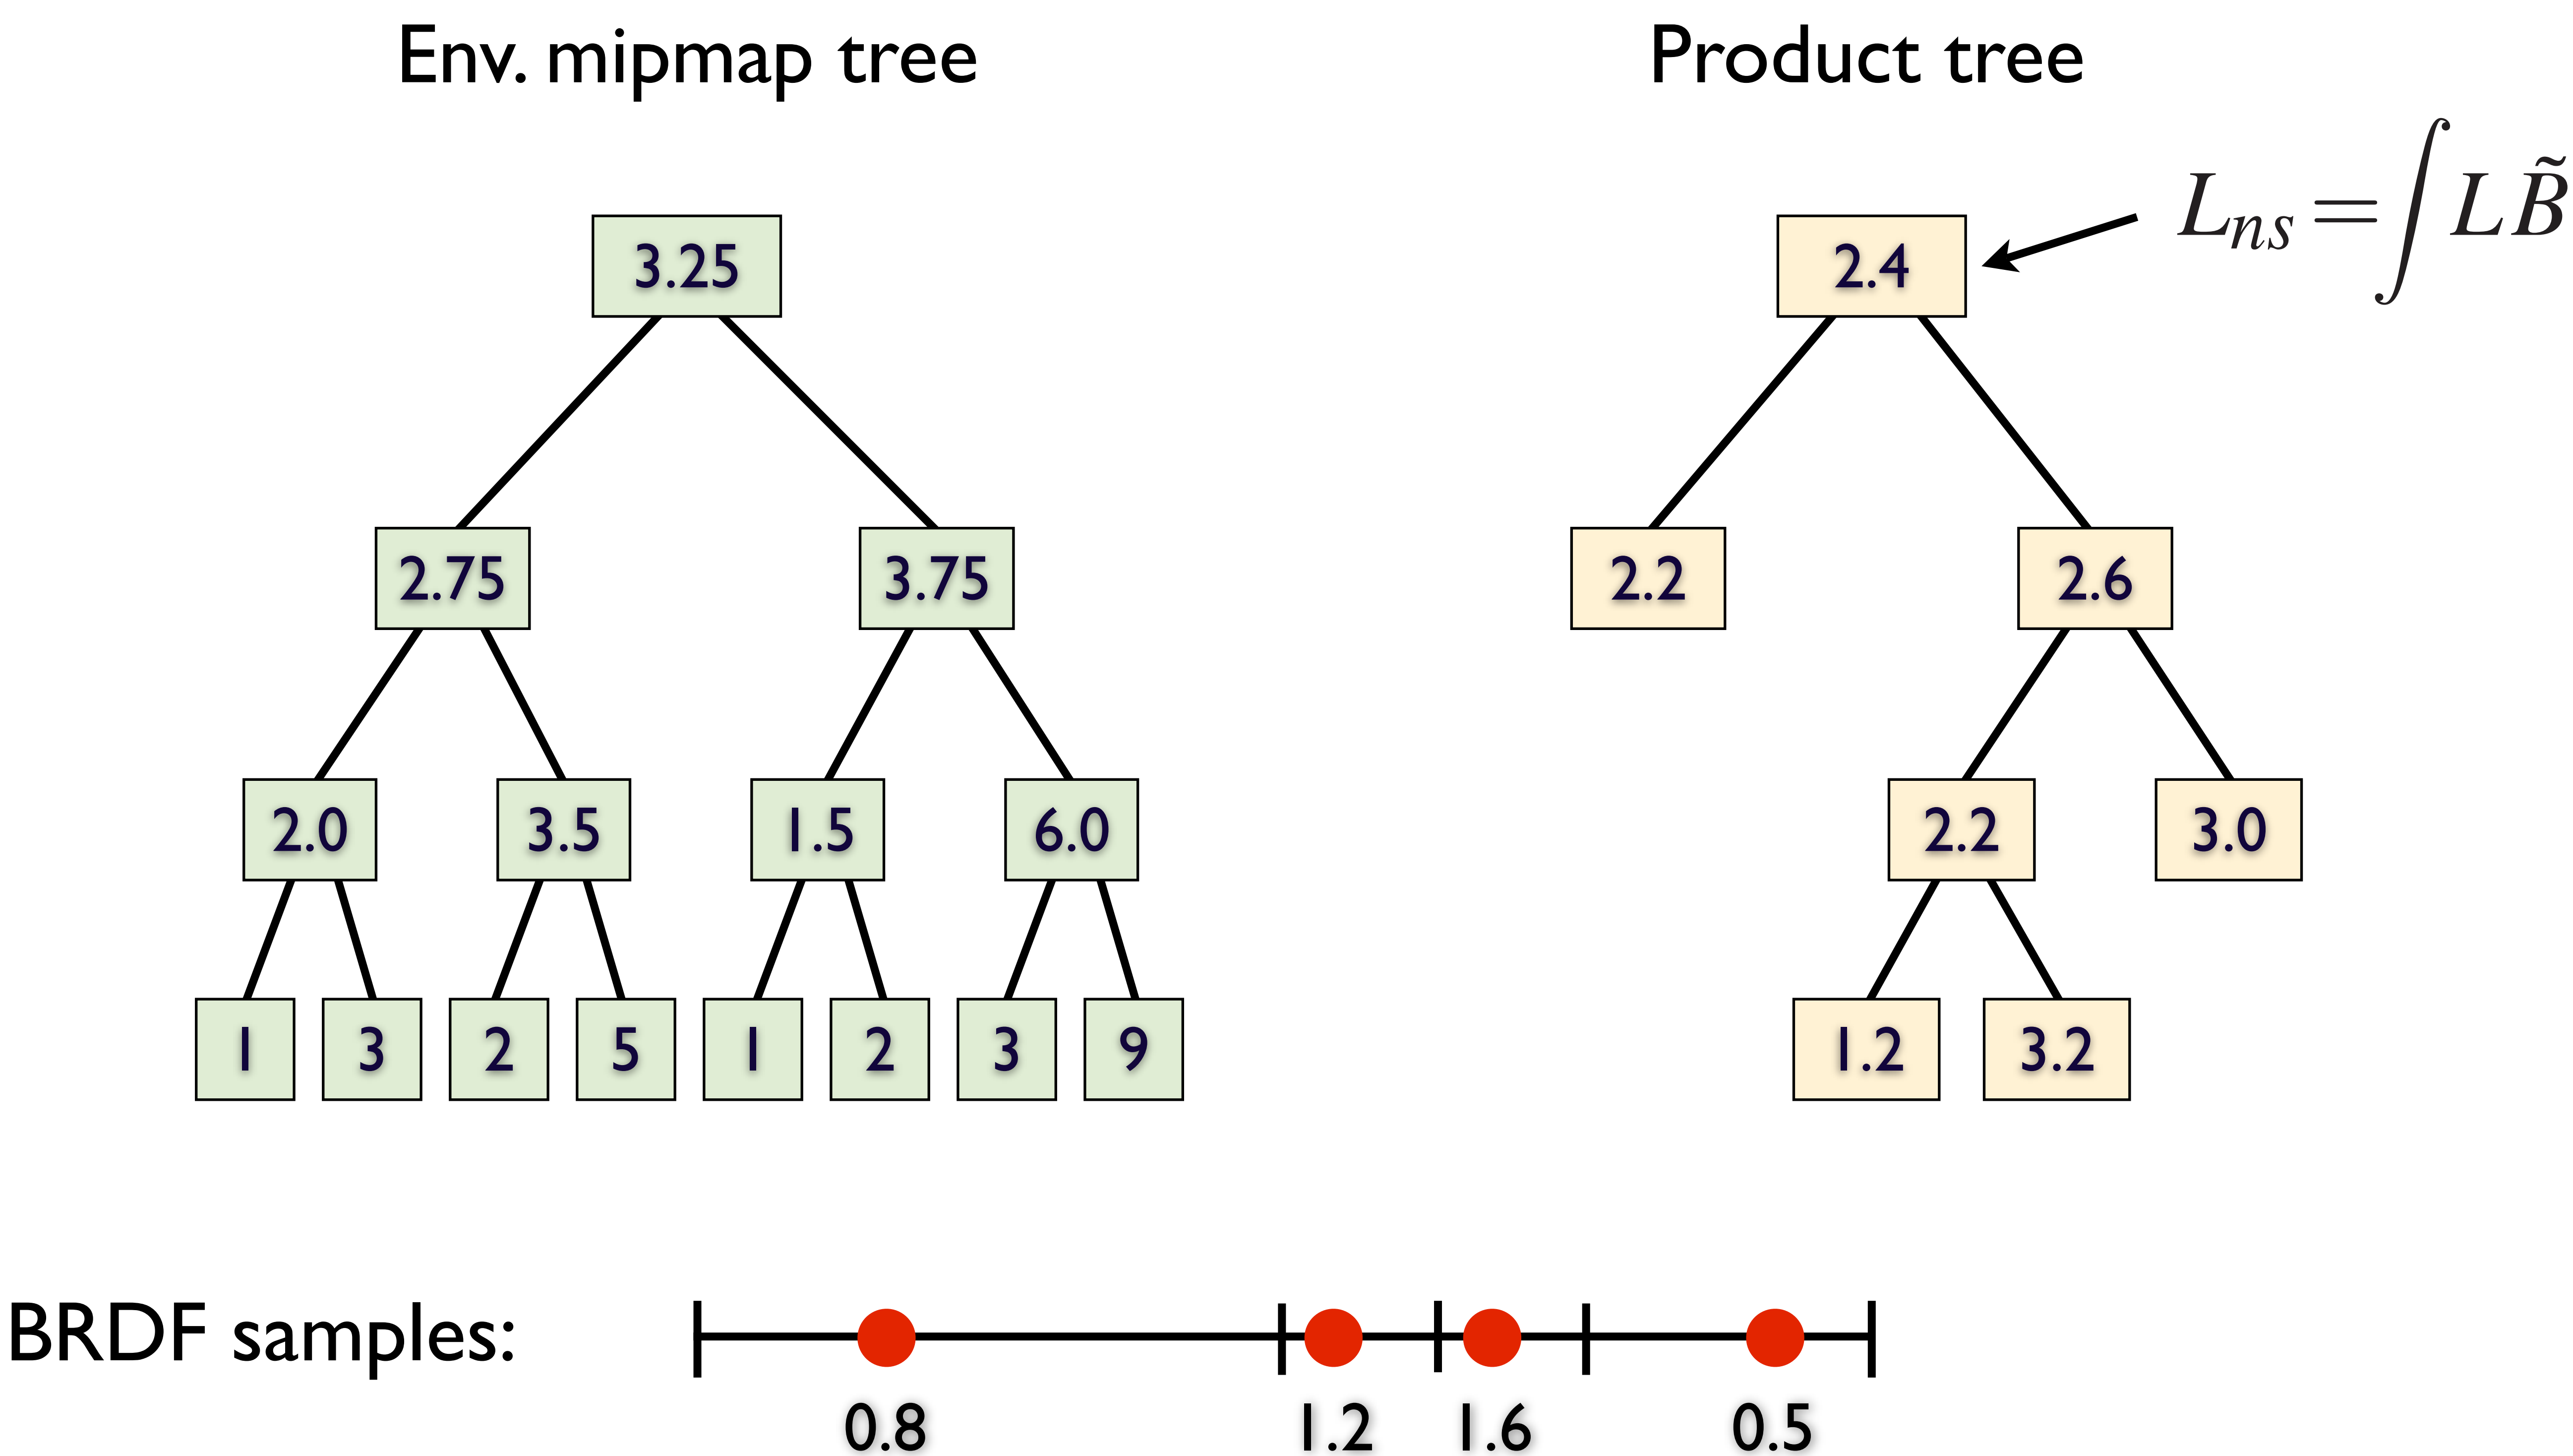

Practical Product Importance Sampling for Direct Illumination, Petrik Clarberg and Tomas Akenine-Möller, EG 2008

The quadtree product is obtained by hierarchically subdividing the sample set until only one sample per node remains. These leaf nodes are assigned the product of the sample value and the corresponding node in the environment map hierarchy. The leaf values are averaged and propagated up towards the root, in a single tree traversal. Pseudo-code can be found in the paper.

Product Sampling

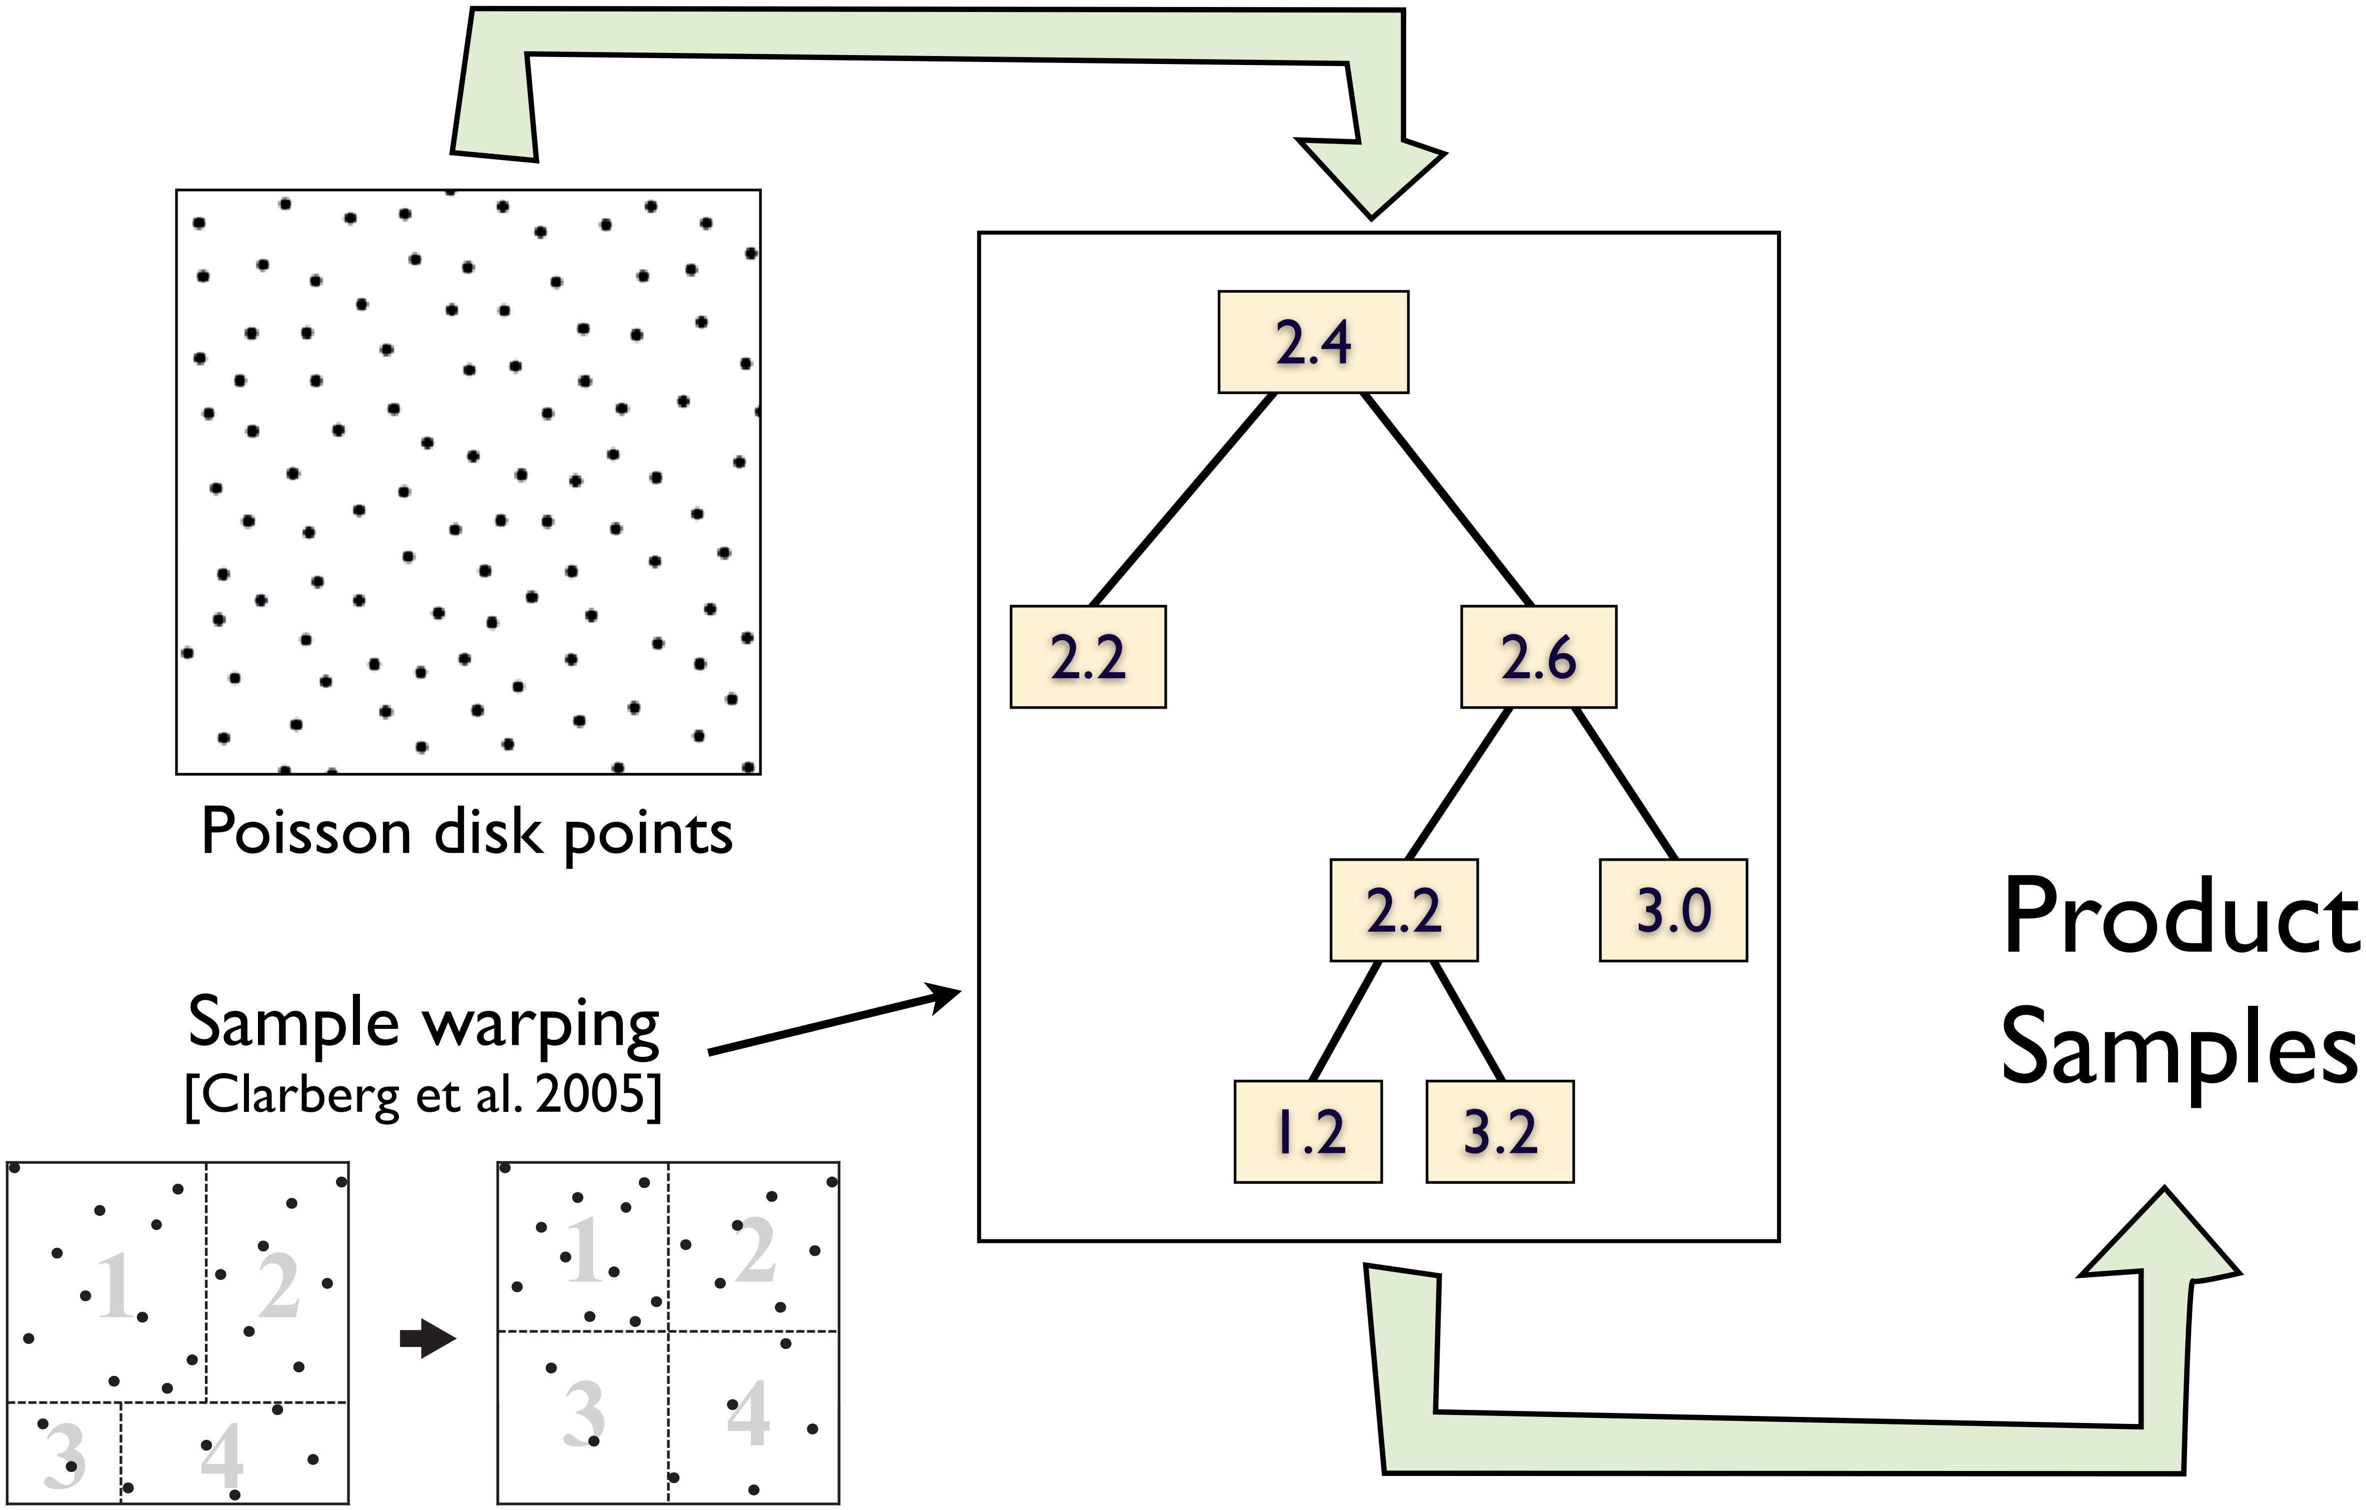

Practical Product Importance Sampling for Direct Illumination, Petrik Clarberg and Tomas Akenine-Möller, EG 2008

Once the quadtree product has been computed, it is sampled using sample warping [Clarberg et al. 2005]. A uniform point set (Poisson disk points) is fed through the tree, and the result is a set of points with the correct product distribution. Pseudo-code and more details can be found in the paper.

Sources of Noise

Our importance function: $$p(\omega) = \frac{L(\omega)\tilde{B}(\omega)}{L_{ns}}$$

- Exact lighting
- Approx. reflectance (BRDF x cosine)

MC estimator: $$L_o = \frac{1}{N}\sum\frac{LBV}{p(\omega)} = \frac{L_{ns}}{N}\sum\frac{B}{\tilde{B}}V$$

Noise

- Visibility
- BRDF approx. inaccuracy
- Tradeoff between #samples for BRDF and visibility

Practical Product Importance Sampling for Direct Illumination, Petrik Clarberg and Tomas Akenine-Möller, EG 2008

Since we use the exact full-res lighting, L is removed from the Monte Carlo estimator. Hence, all noise comes from the visibility (evaluated through ray tracing) and the inaccuracy of the BRDF approximation. L_ns represent the unoccluded lighting and is given by the root node of the product quadtree.

Results

BRDF vs. Visibility Samples (equal time)

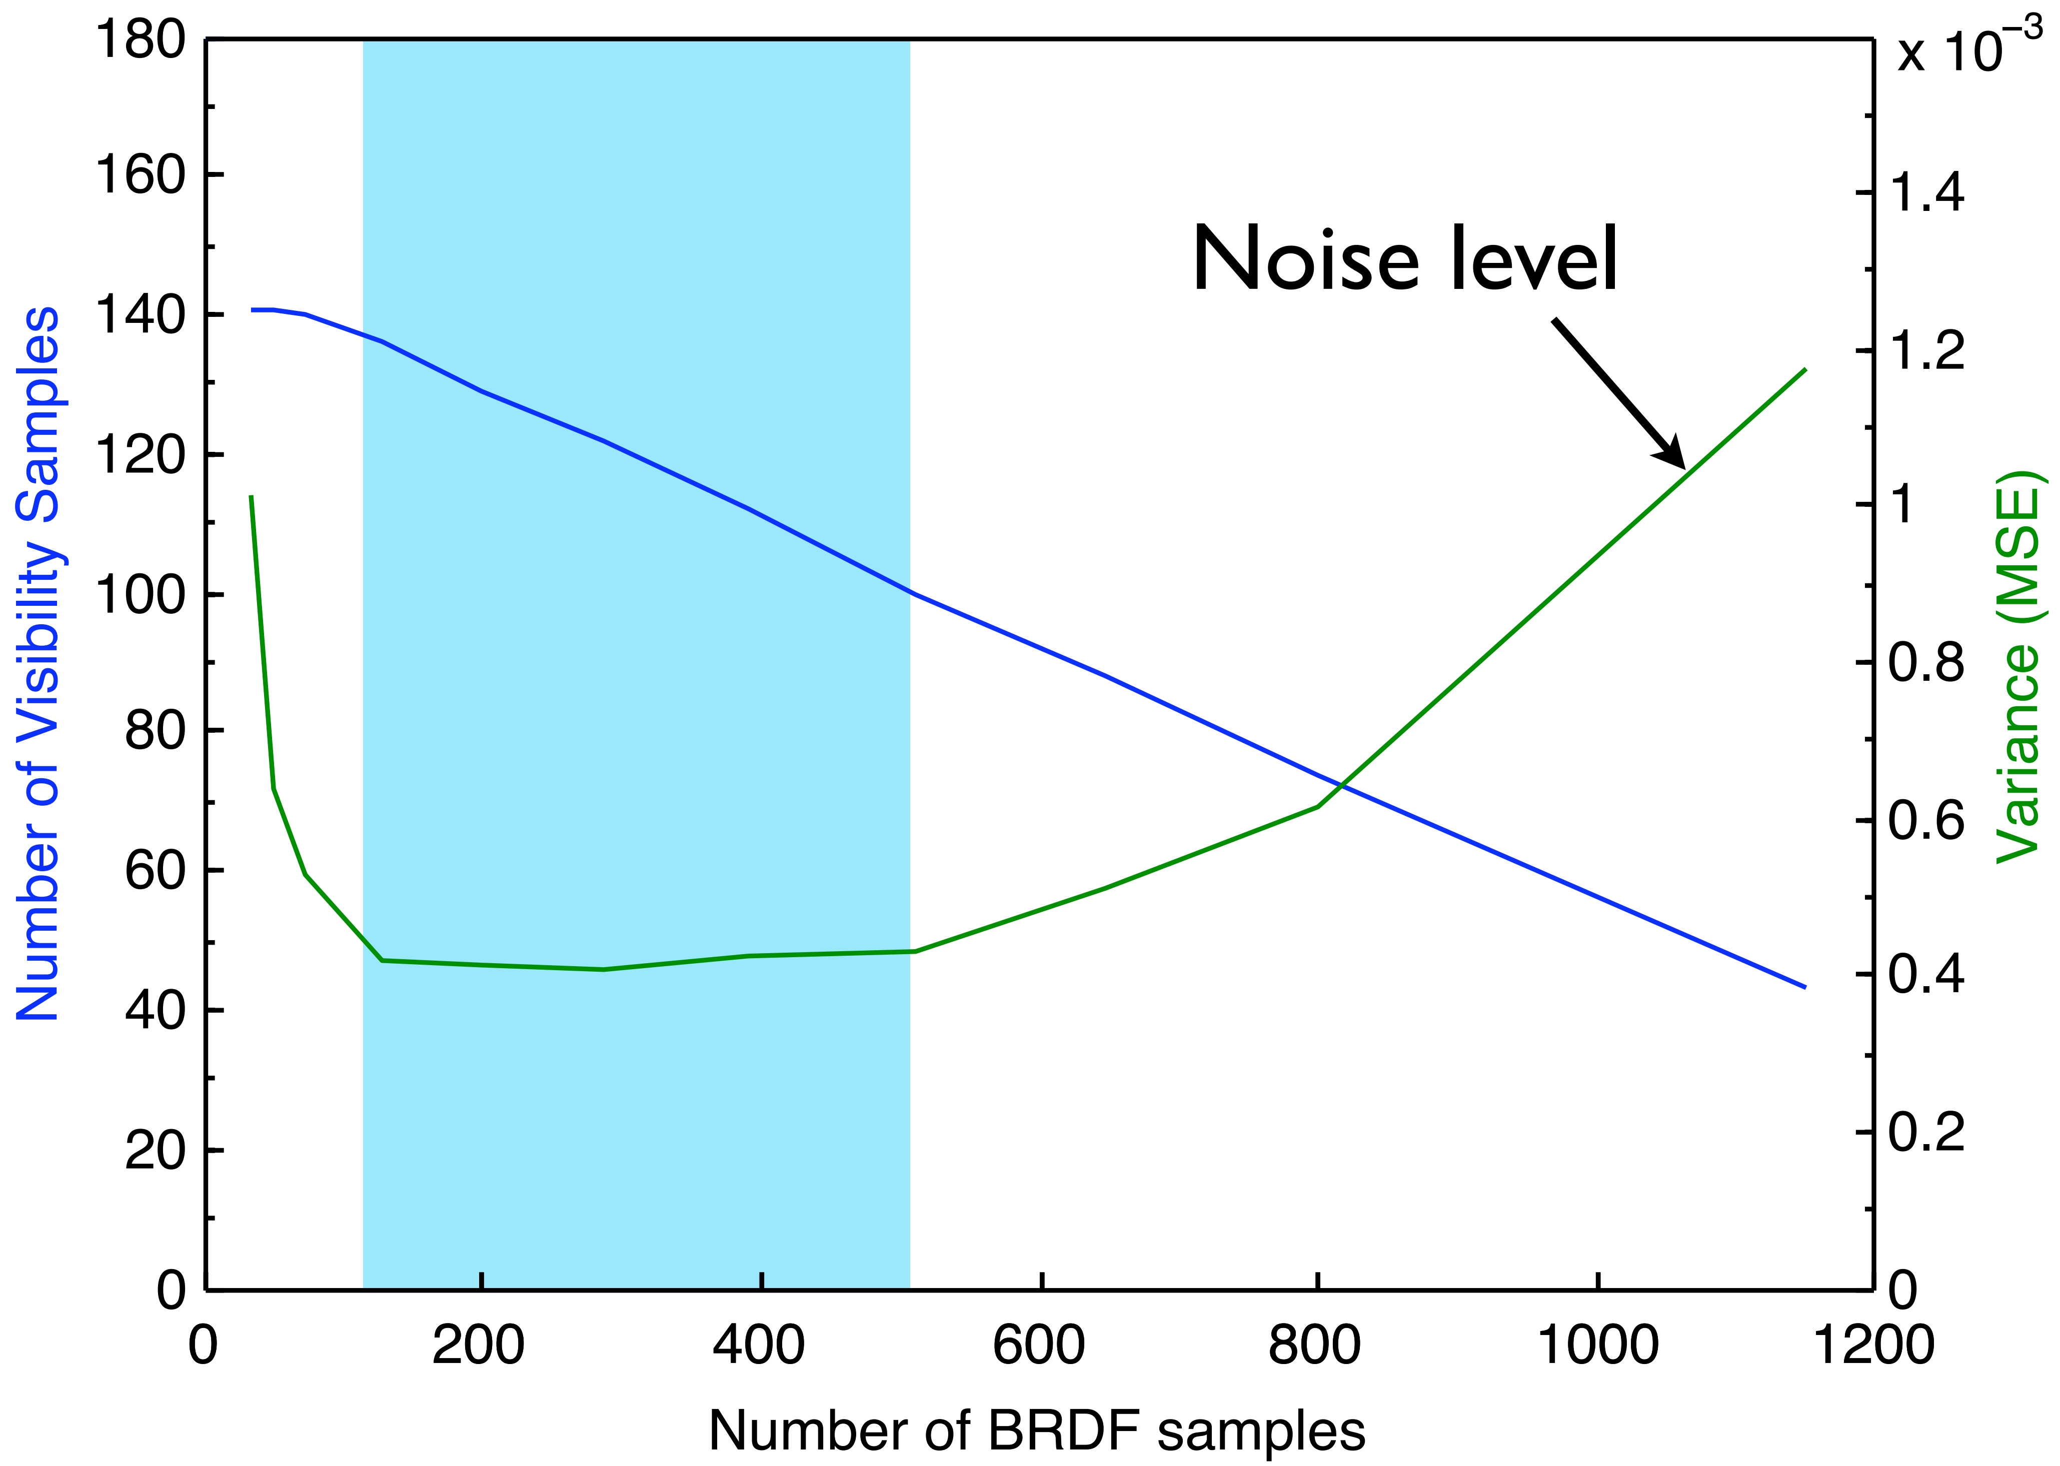

Practical Product Importance Sampling for Direct Illumination, Petrik Clarberg and Tomas Akenine-Möller, EG 2008

First we analyze the tradeoff between the number of visibility samples vs. number of BRDF samples. The variance is low over a quite large range (100–500 BRDF samples), which shows that this is not a critical parameter for our method. A default value of, e.g., 256 BRDF samples works well in most cases.

Test Scene

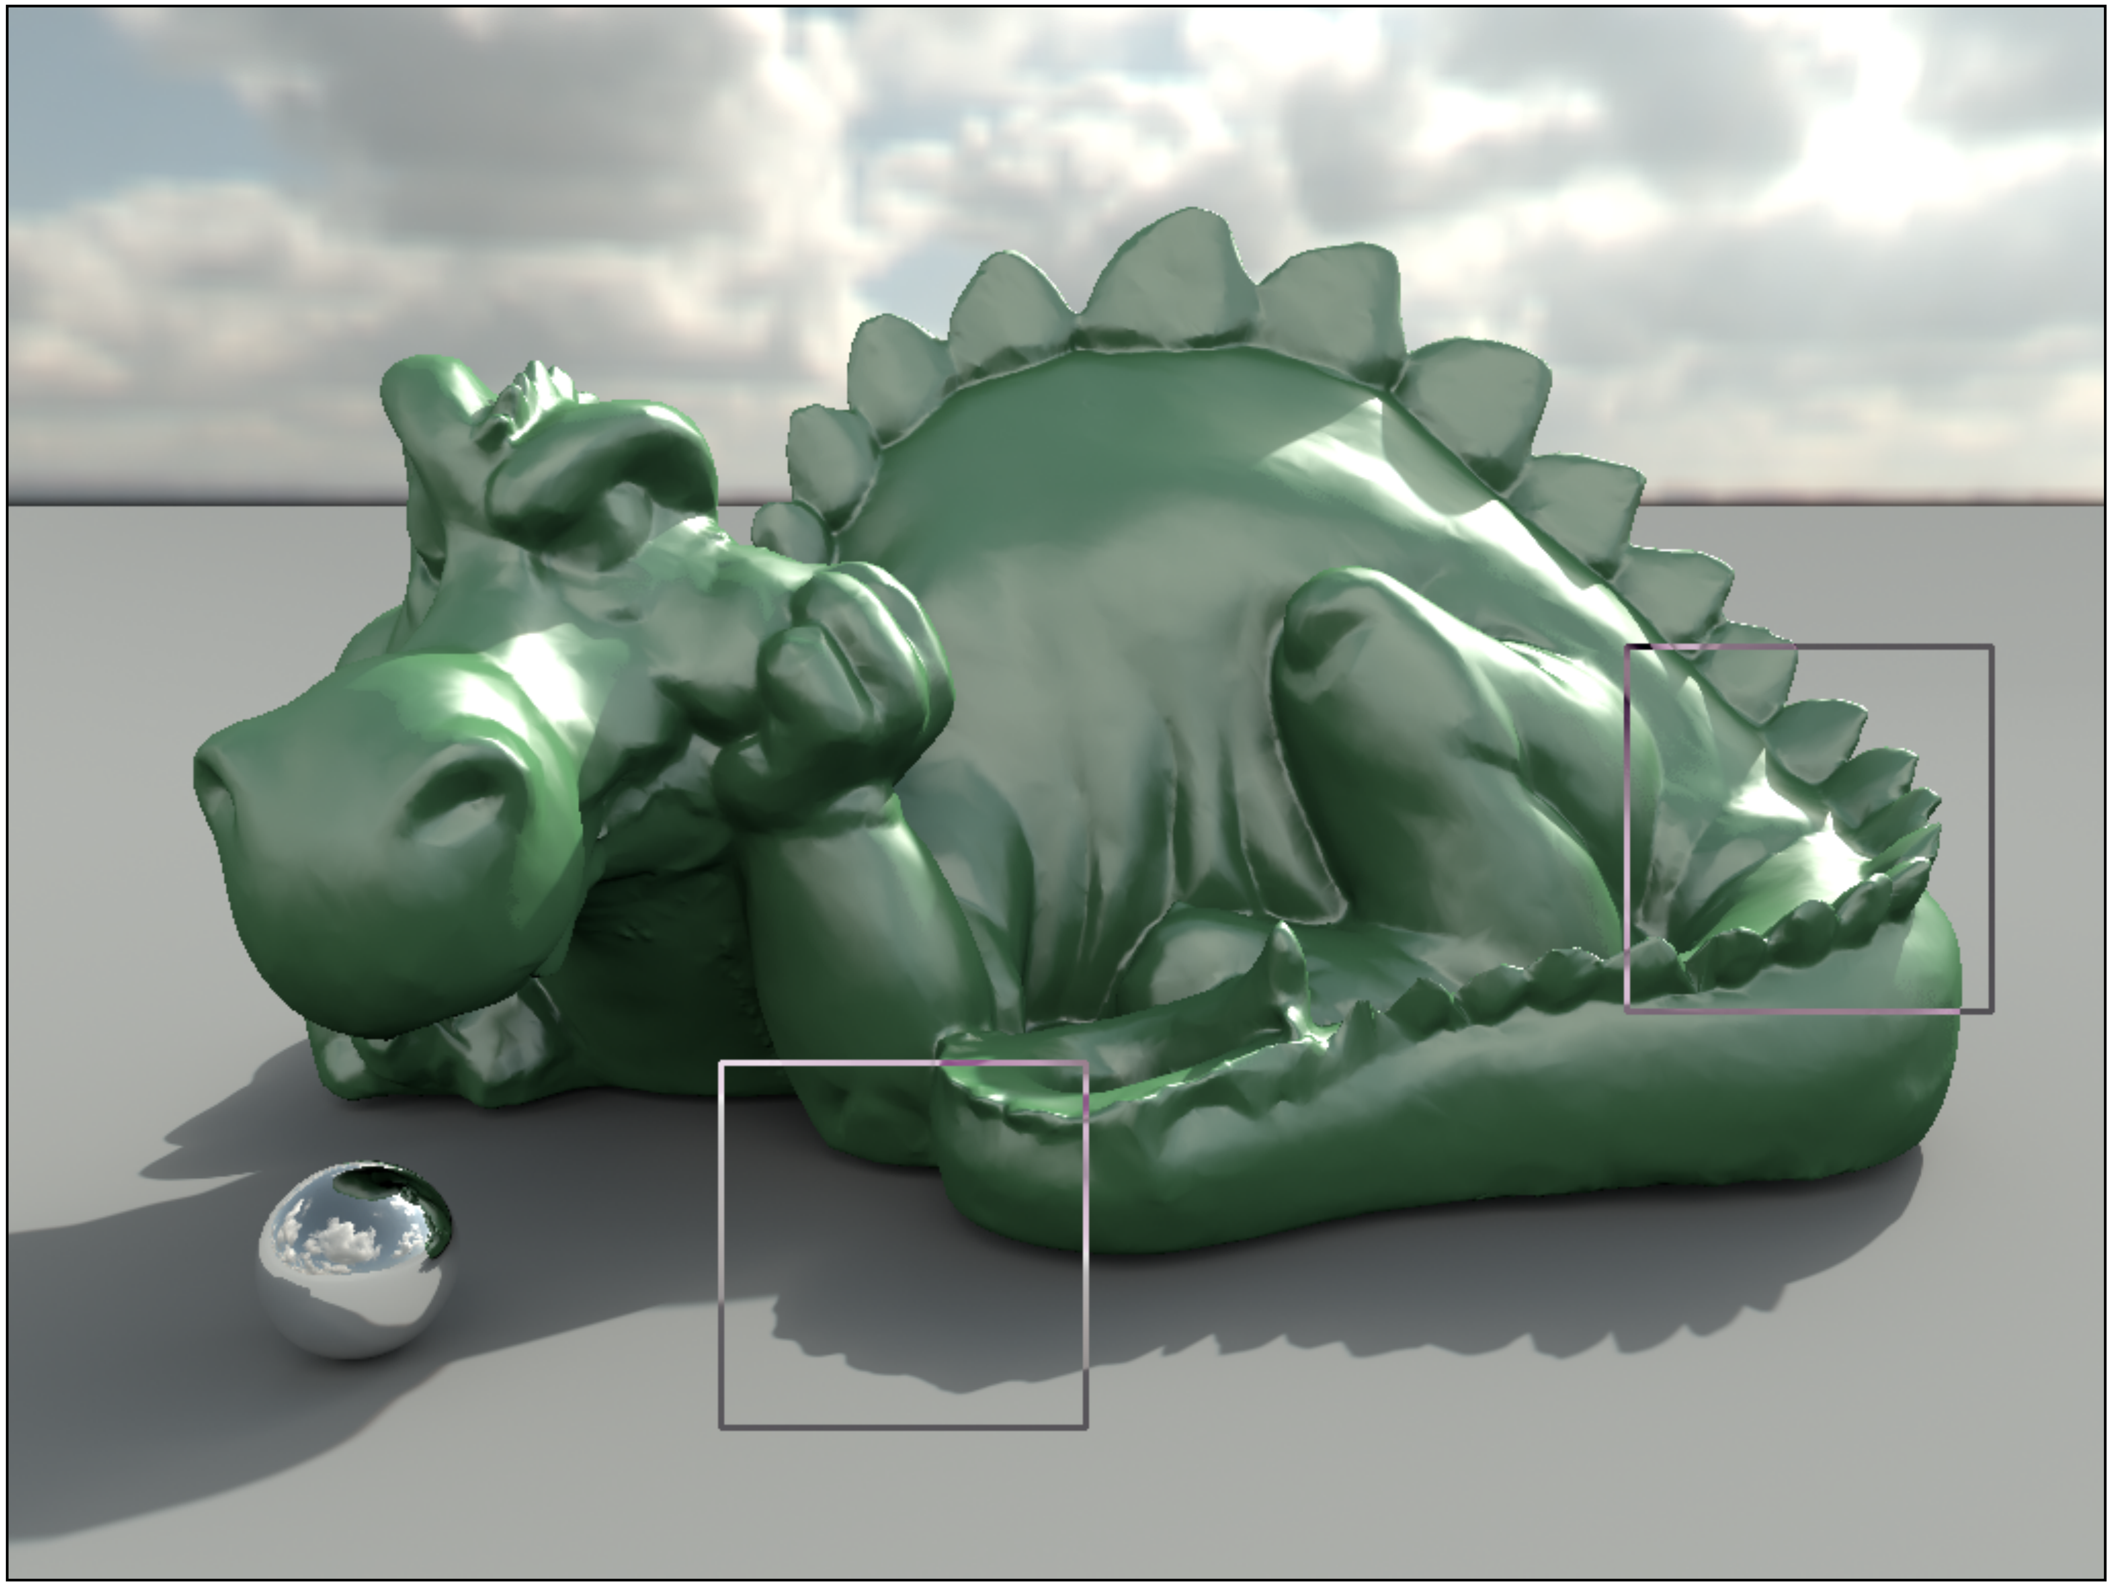

Practical Product Importance Sampling for Direct Illumination, Petrik Clarberg and Tomas Akenine-Möller, EG 2008

Our test scene rendered with an environment map featuring both a single small strong light source (the sun) and larger bright regions (the clouds and sky).

10 samples/pixel

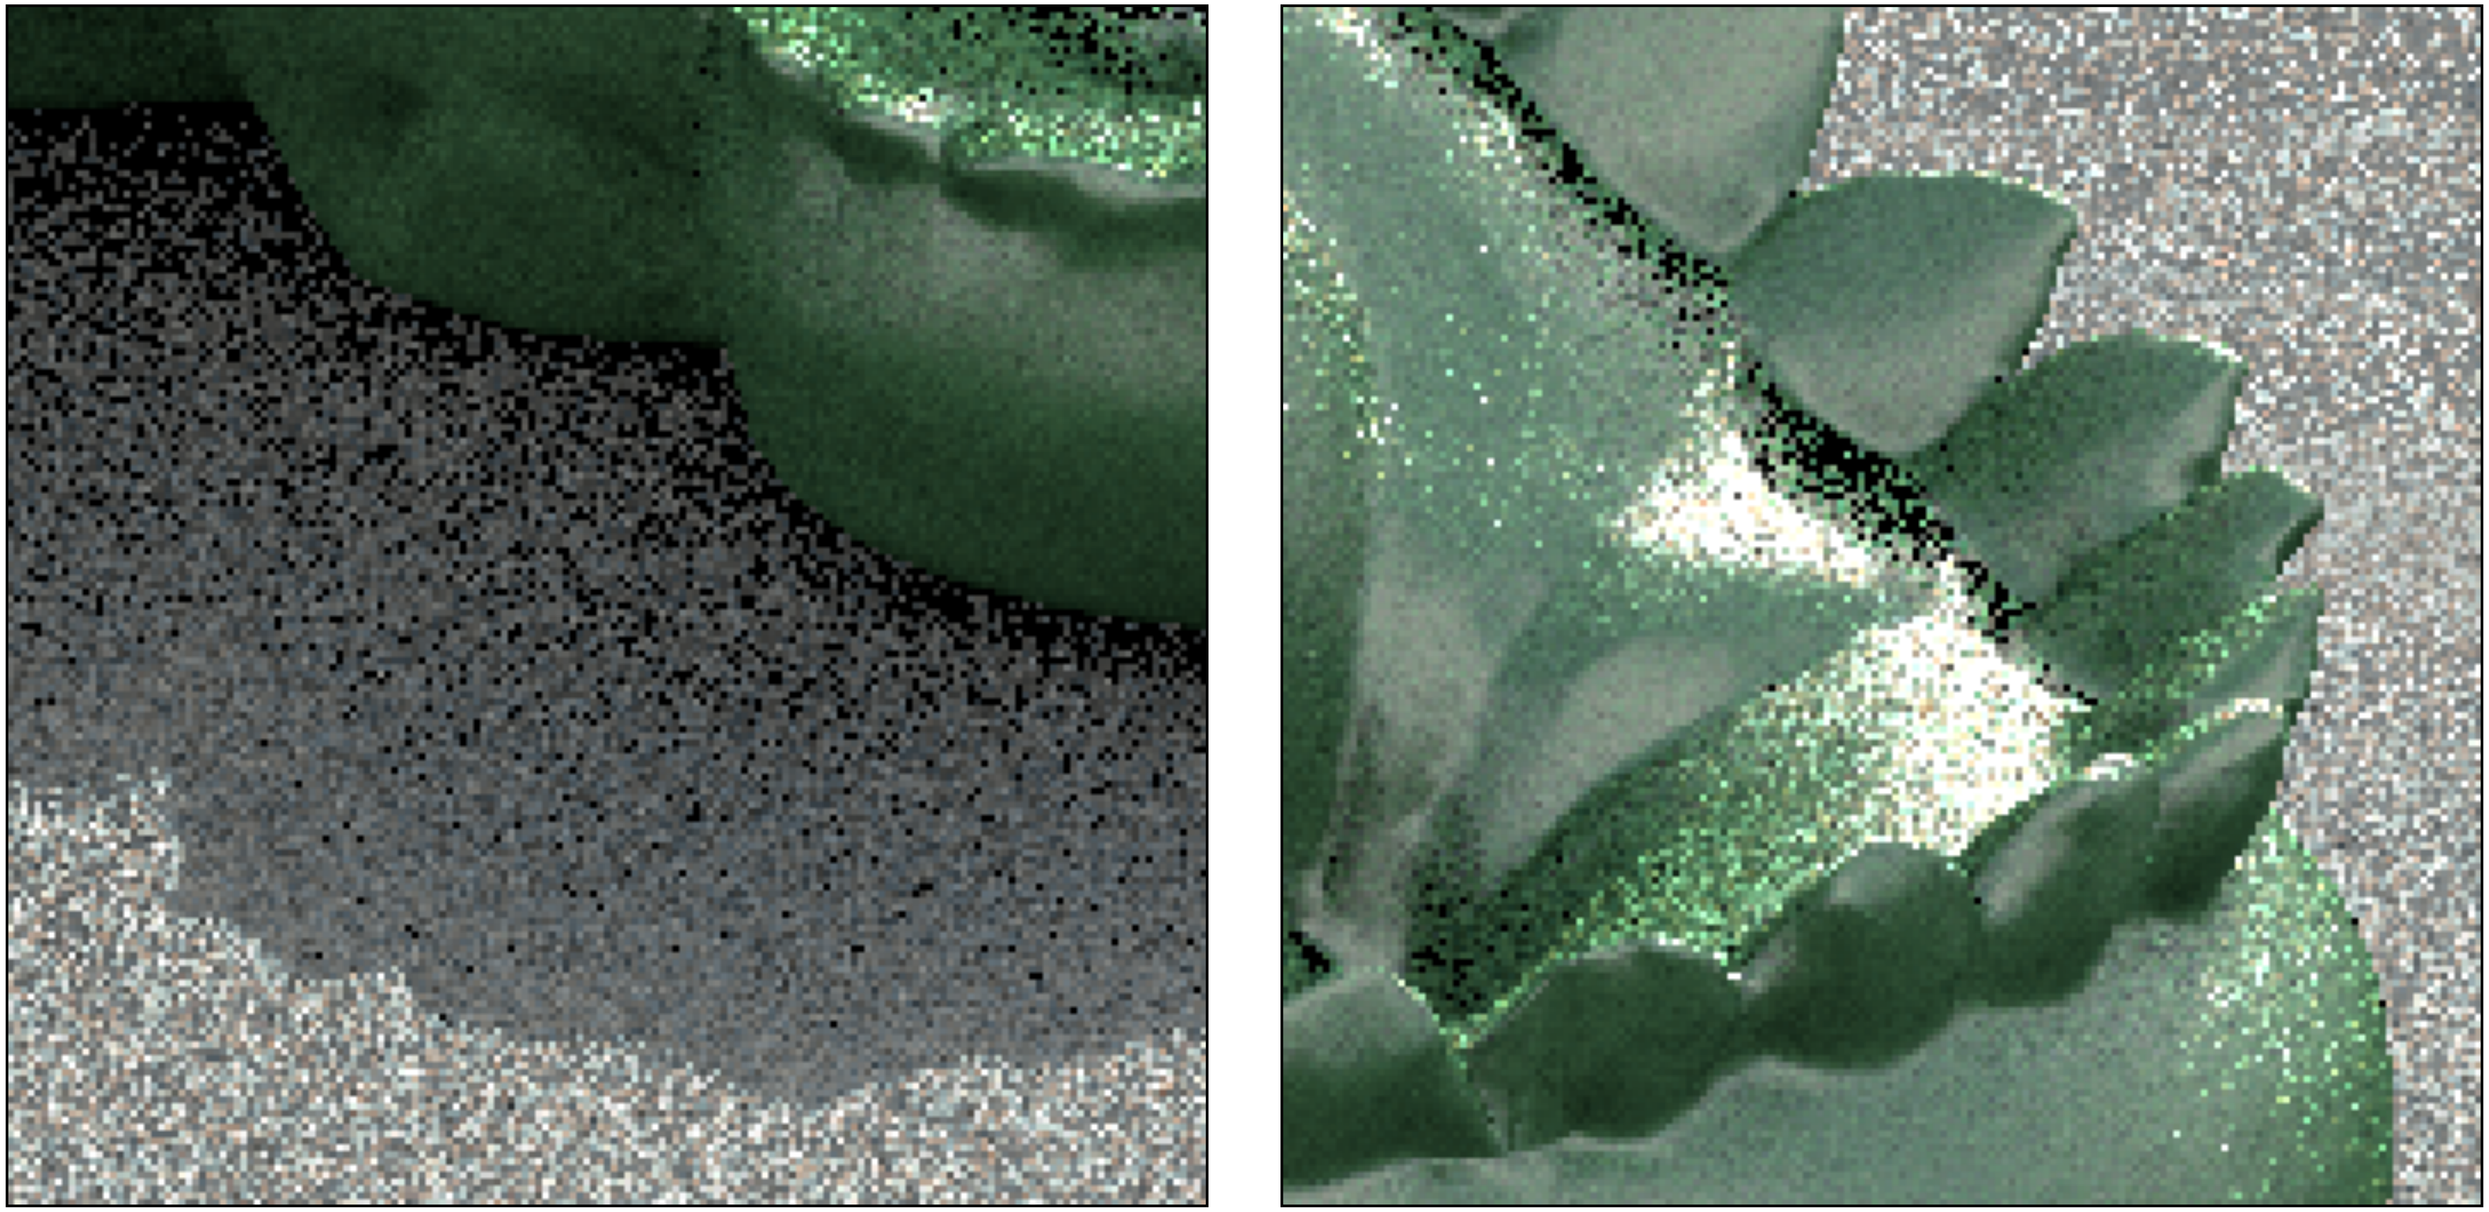

[Clarberg et al. 2005]

10 samples/pixel

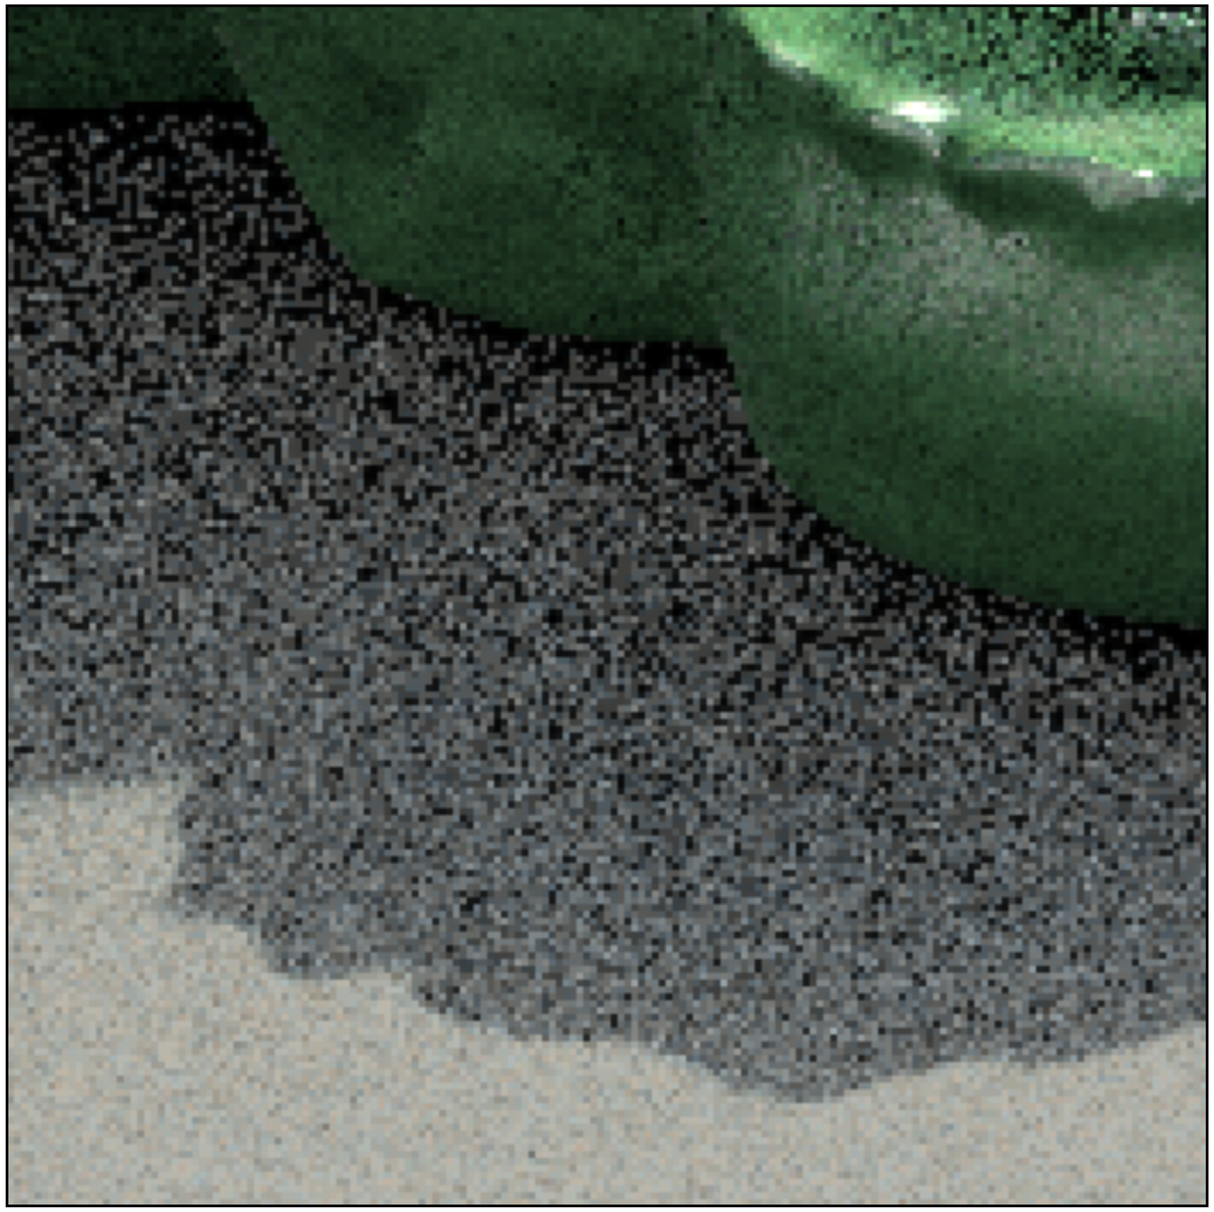

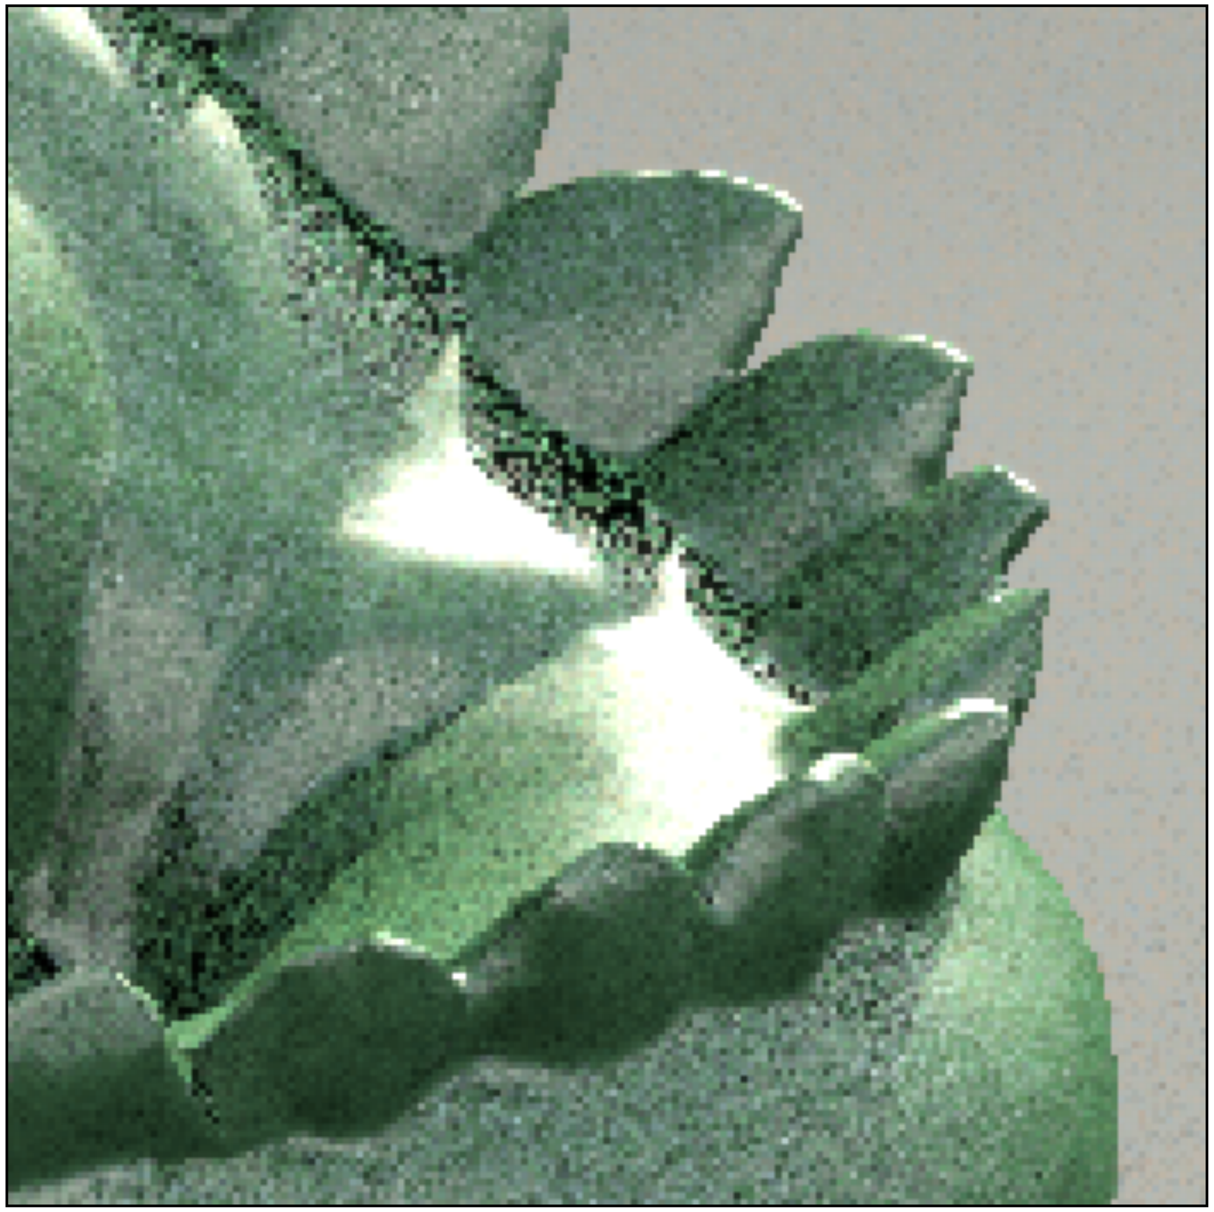

[Cline et al. 2006]

10 samples/pixel

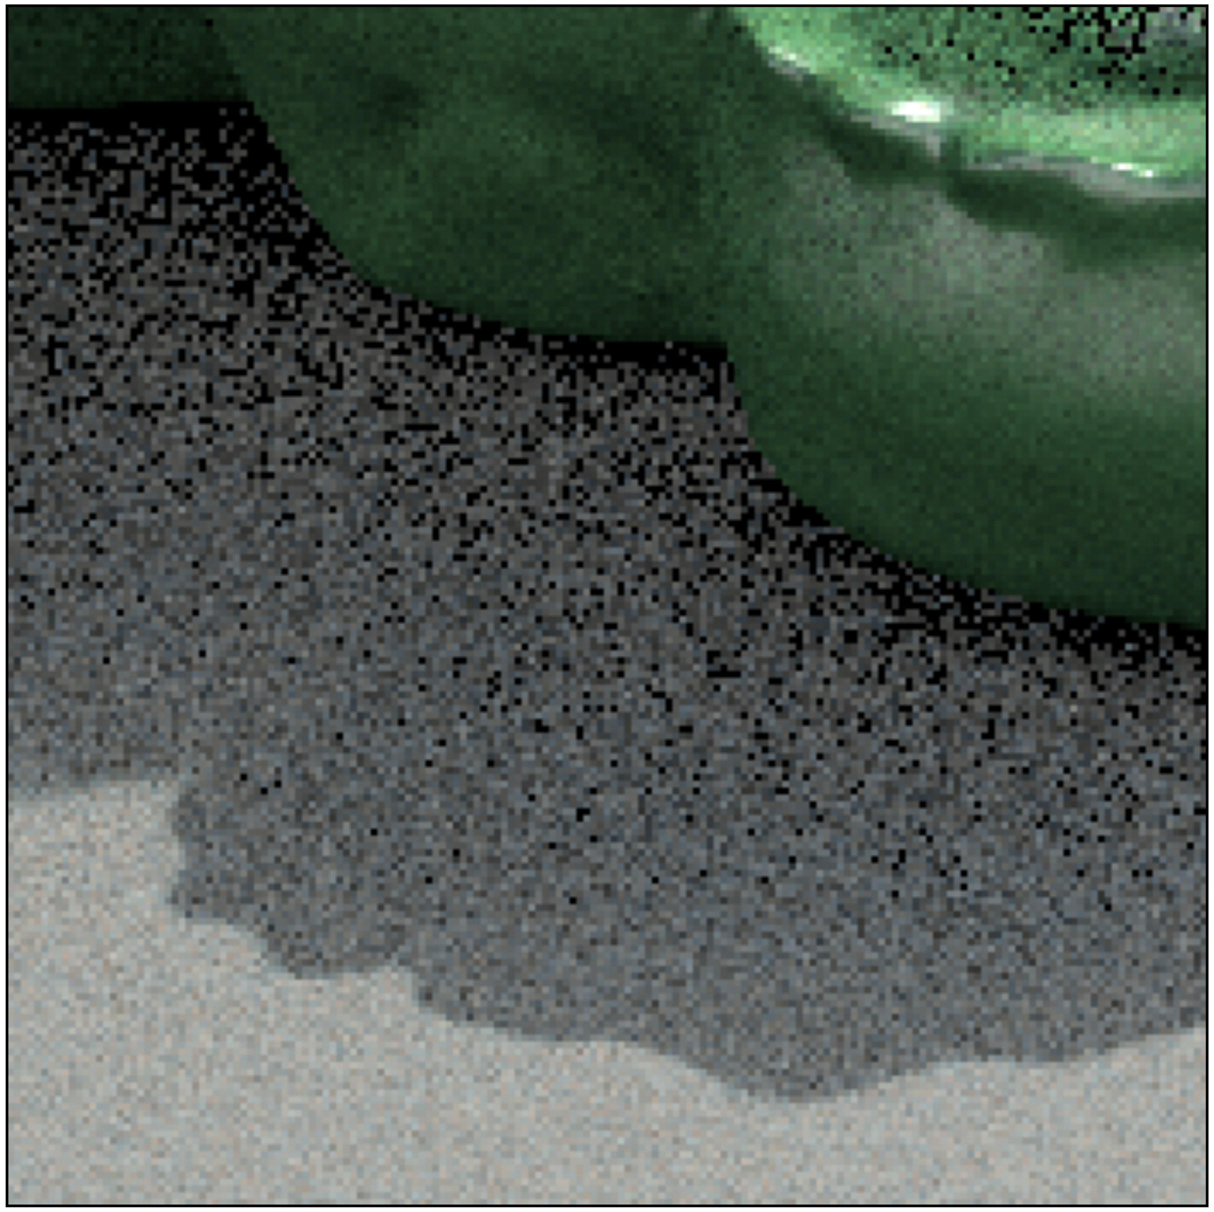
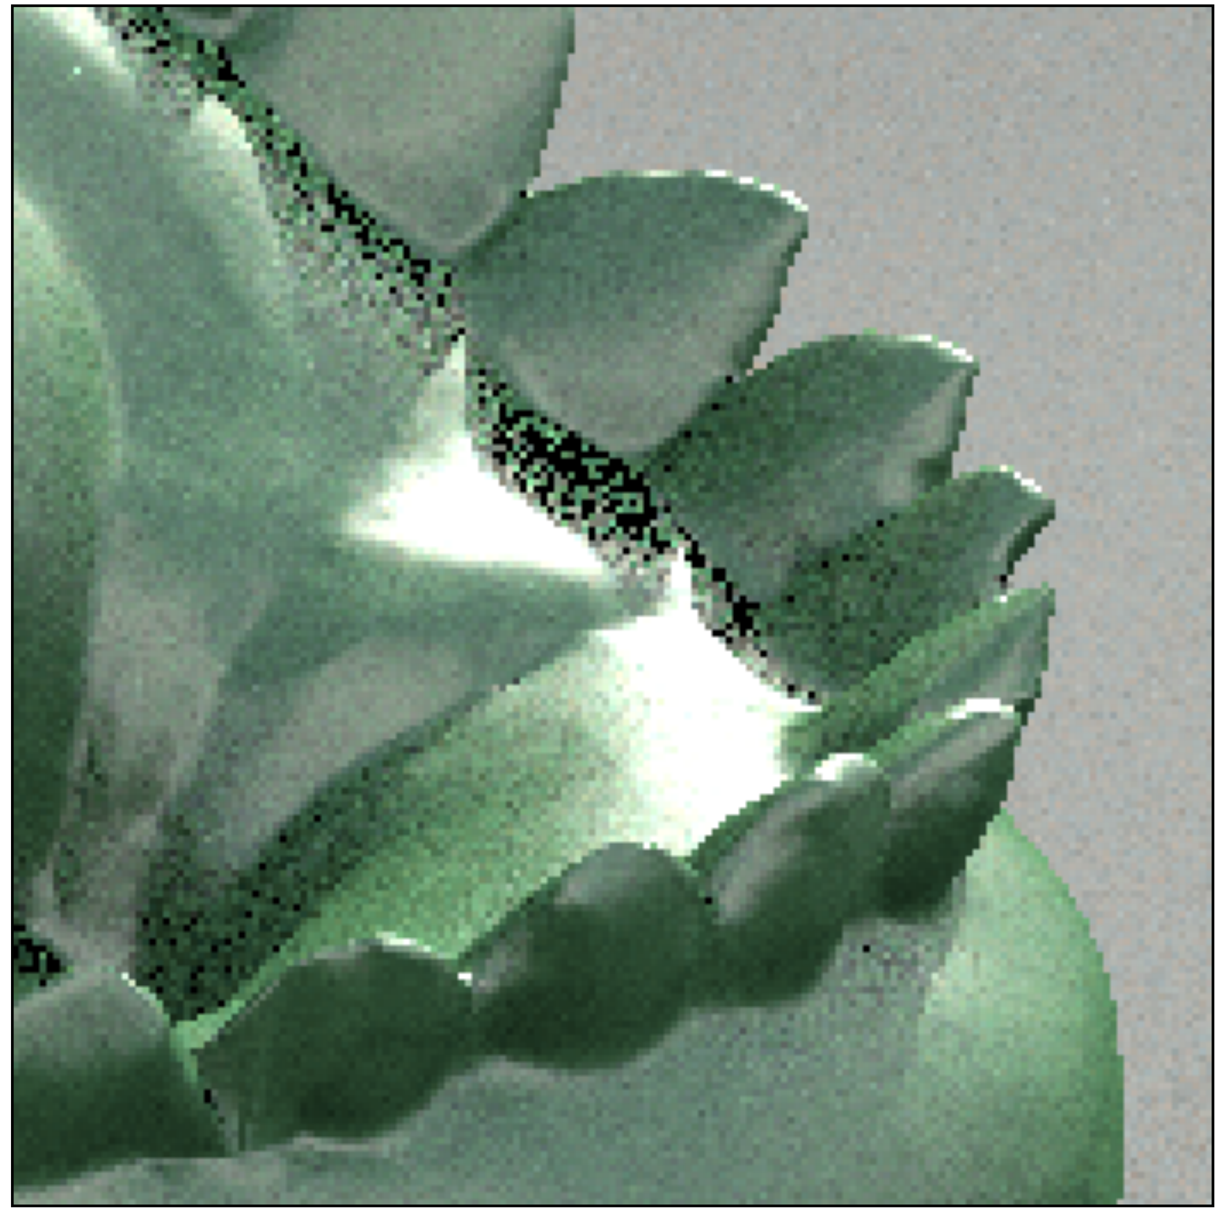

Our algorithm

Practical Product Importance Sampling for Direct Illumination, Petrik Clarberg and Tomas Akenine-Möller, EG 2008

Rendering Time (seconds)

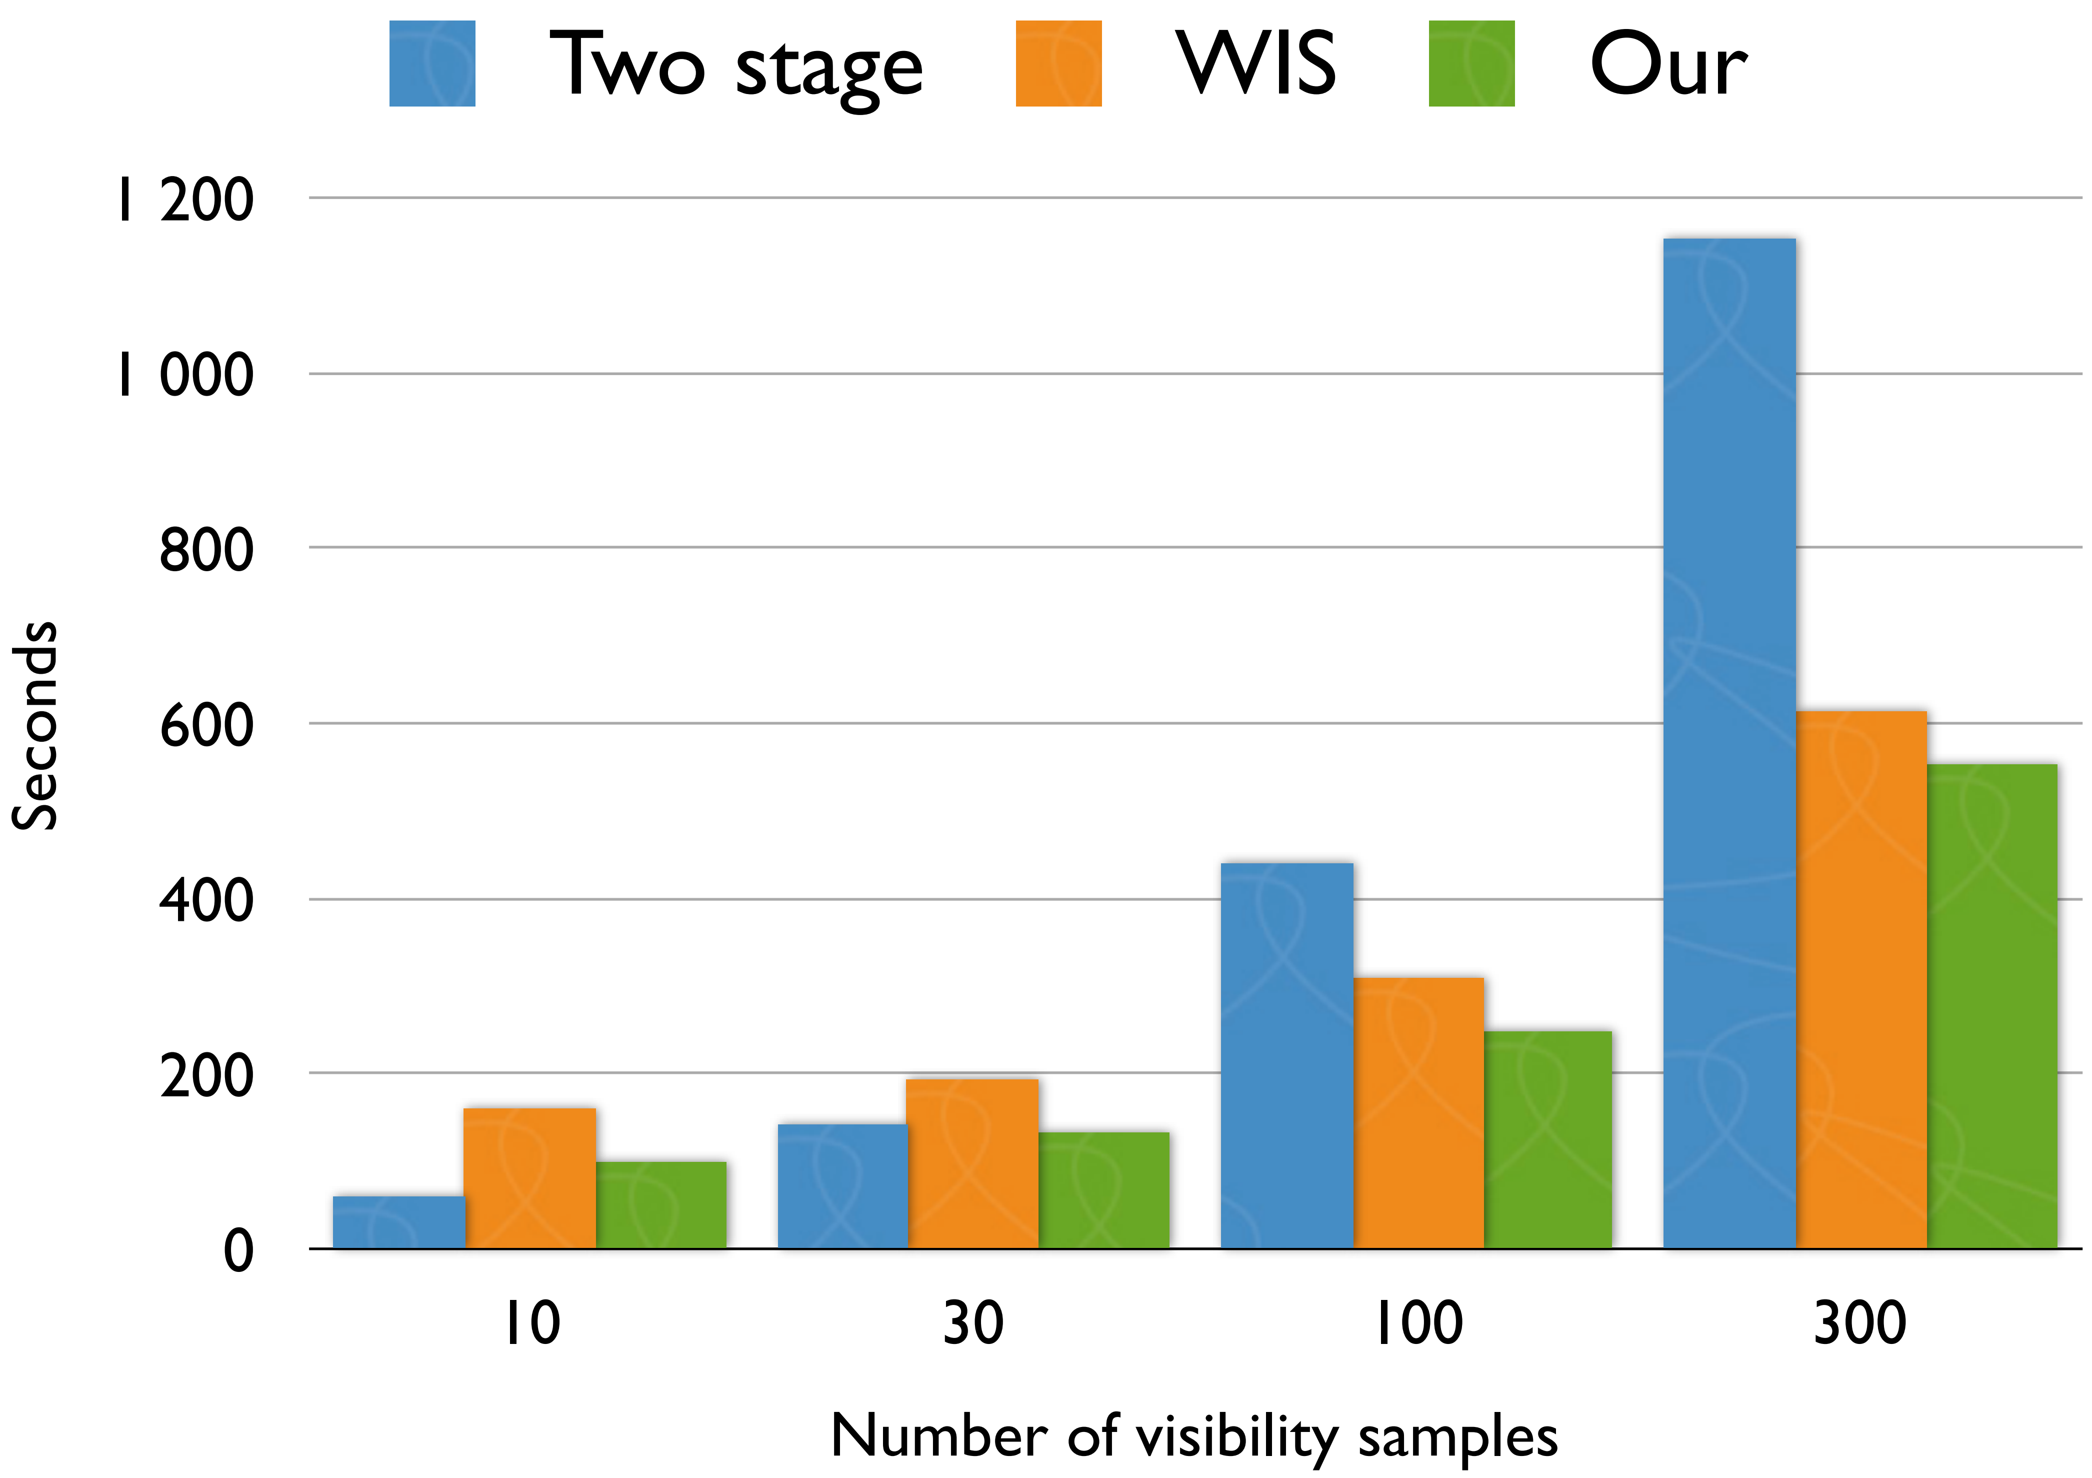

Practical Product Importance Sampling for Direct Illumination, Petrik Clarberg and Tomas Akenine-Möller, EG 2008

Two stage importance sampling [Cline et al. 2006] is fast at low sample rates, but gets slower with many samples. Between our method and WIS [Clarberg et al. 2005], there is a constant small speedup, which comes from the simpler evaluation of the product (quadtree vs. wavelets). Note that our method also performs BRDF sampling on-the-fly, while WIS used precomputed BRDFs.

Equal-time Comparison

[Cline et al. 2006]

Our algorithm

Practical Product Importance Sampling for Direct Illumination, Petrik Clarberg and Tomas Akenine-Möller, EG 2008

Side-by-side comparison between Two stage importance sampling and our method. Equal-rendering time and about 30 samples/pixel.

Equal-time Comparison

[Cline et al. 2006]

Equal-time Comparison

Our algorithm

Contributions

Picked the best from previous algorithms:

- No precomputation
- Little memory use
- Flexible BRDF support
- High resolution lighting
- Simple implementation

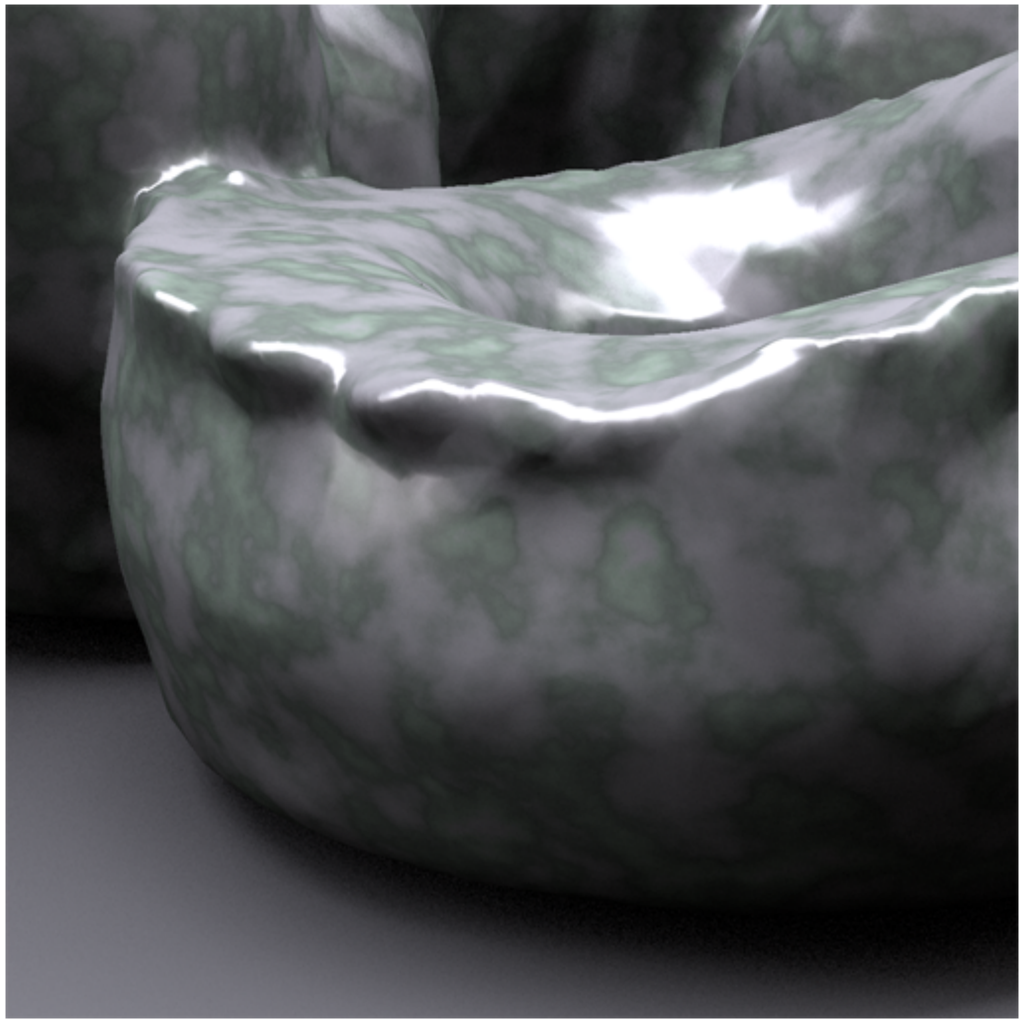

Fast quadtree product and wavelet product

Acknowledgements

Funding

- Swedish Foundation for Strategic Research
- Intel Corporation

Lund University Graphics Group

- Jacob Munkberg, Jon Hasselgren, and the rest

The End