



Wolfgang Tatzgern, BSc

# **Stochastic Substitute Trees for Real-Time Global Illumination**

## **Master's Thesis**

to achieve the university degree of

Master of Science

Master's degree programme: Computer Science

submitted to

**Graz University of Technology**

Supervisor

Ass.Prof. Dipl.-Ing. Dr.techn. BSc Markus Steinberger

Institute of Computer Graphics and Vision

Head: Univ.-Prof. Dipl.-Ing. Dr.techn. Dieter Schmalstieg

Graz, January 2020

This document is set in Palatino, compiled with [pdfL<sup>A</sup>T<sub>E</sub>X2<sub>ε</sub>](#) and [Biber](#).

The L<sup>A</sup>T<sub>E</sub>X template from Karl Voit is based on [KOMA script](#) and can be found online: <https://github.com/novoid/LaTeX-KOMA-template>

## Affidavit

I declare that I have authored this thesis independently, that I have not used other than the declared sources/resources, and that I have explicitly indicated all material which has been quoted either literally or by content from the sources used. The text document uploaded to TUGRAZonline is identical to the present master's thesis.

---

Date

---

Signature



# Acknowledgement

First I would like to thank my supervisor Markus Steinberger for the continuous support, guidance and advice he has provided throughout this project. I would also like to give a big thank you to Bernhard Kerbl for all of his input, support and encouragement.

Also, I thank Benedikt Mayr for providing a solid foundation for this work. Furthermore, I would like to thank the Faculty of Computer Science and Biomedical Engineering from Graz University of Technology for funding this work through the Promotional Scholarship.

Last but not least, I want to thank my beloved parents Christine and Josef and of course Karina for more reasons than I could ever possibly list.



# Abstract

With the introduction of hardware-supported raytracing and deep-learning for denoising, computer graphics has made a considerable step toward real-time global illumination. However, the per-pixel ray budget is still too low to support extensive path tracing. In this work, we present an alternative global illumination method: The stochastic substitute tree (SST), a hierarchical structure inspired by lightcuts with light probability distributions as inner nodes. Our approach distributes virtual point lights (VPLs) in every frame using hardware raytracing and efficiently constructs the SST over those lights by clustering according to Morton codes. Rebuilding the SST in every frame allows us to work with a moderate number of VPLs as their locations change rapidly from frame-to-frame while enabling dynamic lighting scenarios. Global illumination is approximated by sampling the SST and considers the BRDF at the hit location as well as the SST nodes' intensities for importance sampling directly from inner nodes of the tree. To remove the introduced Monte Carlo noise, we use a pruned recurrent autoencoder. In combination with temporal filtering, we deliver real-time global illumination for complex scenes with challenging light distributions.





# Contents

<b>Abstract</b>	<b>vii</b>
<b>1 Introduction</b>	<b>1</b>
<b>2 Background and Related Work</b>	<b>3</b>
<b>3 Stochastic Substitute Trees</b>	<b>7</b>
3.1 The Substitute Tree . . . . .	7
3.2 Distribution Merging and Errors . . . . .	10
3.3 Parallel Construction . . . . .	11
<b>4 Stochastic Light Sampling</b>	<b>15</b>
4.1 Tree traversal . . . . .	15
4.2 Sampling the Substitute Distribution . . . . .	17
<b>5 Implementation</b>	<b>19</b>
<b>6 Evaluation</b>	<b>21</b>
6.1 Comparison . . . . .	23
6.2 Scalability . . . . .	26
6.3 Discussion and Limitations . . . . .	28
<b>7 Conclusion</b>	<b>31</b>
<b>Bibliography</b>	<b>33</b>



# 1 Introduction

Real-time global illumination (GI) is one of the most sought-after features in computer graphics. Path tracing renownedly provides an elegant solution to simulating global light transport by recursively sampling along a large number of rays that traverse the scene. In the film industry, path tracing has become the method of choice and is usually performed off-line on the central processing unit (CPU). However, its image-order complexity prevents it from being used in the same way for driving real-time applications. Today, almost all competitive path tracers leverage the massively parallel processing power of the graphics processing unit (GPU). Recent GPUs have been fitted with dedicated ray-tracing modules that help mitigate the bottleneck of ray-based scene traversal. Furthermore, specialized inferencing units (*tensor cores*) have been added to evaluate pre-trained convolutional neural networks (CNN) to filter undersampled images. Although backed by hardware, ray-tracing and inference remain prohibitively expensive. Hence, these novel GPU features are mostly used in combination with rasterization for bulk geometry rendering, contributing minor visual enhancements.

In contrast to path tracing, *instant radiosity* (IR) describes an alternative solution to GI on hardware rasterization. IR effectively reduces the problem of indirect illumination to trivial direct illumination by introducing a large number of *virtual point lights* (VPL) that are distributed throughout the scene. In addition to avoiding image-order complexity, such *many-light* rendering approaches are particularly effective when combined with the concept of *lightcuts*: instead of full surface shading with all lights in the scene, a lightcut selects a subset of VPLs and amplifies their intensity for shading. However, the generation of high-quality lightcuts is computationally intensive and thus usually implies off-line processing.

In this paper, we describe a new approach for fast, high-quality global illumination on modern GPUs by extending the concept of stochastic lightcuts

## 1 Introduction

(SLC) [Yuk19]. Our *Stochastic Substitute Trees* (SST) use custom substitute lights whose geometric properties are not simply adopted from leaf nodes, but rather purpose-built to stochastically approximate their subhierarchies. In the context of real-time GI, we show how SSTs provide an ideal complement to recently introduced GPU capabilities as outlined in Fig. 3.1: in combination with ray-tracing for visibility, deep learning for image reconstruction and adaptive temporal filtering, SSTs compare favorably to other state-of-the-art solutions in terms of performance and temporal stability. We make the following contributions:

- We introduce a novel method for generating substitutes for inner light tree nodes that model their light contribution as probability distributions. To the best of our knowledge, we are the first to propose using representatives whose geometric properties are disjoint from the initial VPLs.
- We show how SSTs can be generated in parallel on the GPU, enabling real-time animations and the application of up-to-date light sampling in every frame.
- We demonstrate how the SST can be efficiently sampled for Lambertian and specular materials.
- We demonstrate our approach in NVIDIA Falcor and evaluate it against state-of-the-art methods for global illumination on a variety of static and animated scenes.
- We discuss the combination with hardware ray-tracing, inferencing capabilities, and temporal filtering.

## 2 Background and Related Work

Following the original rendering equation [Kaj86], path tracing emerged as a reliable Monte Carlo method for global illumination. Given an input scene, recursive, probabilistic sampling of light paths is combined with Monte Carlo integration to yield an initially approximate result that eventually converges to the solution of the rendering equation. To reduce the computational effort required to reach convergence importance sampling, next-event estimation and bi-directional path tracing have been proposed. Programmable GPUs further enabled the implementation of high-performance, massively parallel path tracing variants, leading to an additional boost in their efficacy [Par+10]. While path tracing entire frames remains unattainable in real-time, ray tracing is commonly employed to generate individual effects such as shadows, ambient occlusion or specular reflections [Kel+19]. Although ray tracing is classified as “embarrassingly parallel”, achieving optimal utilization of the GPU has proven challenging [AL09]. State-of-the-art GPU models are therefore fitted with dedicated ray tracing hardware [NVI18; HA19].

Aside from path tracing, the most prominent alternatives for simulating global illumination include *metropolis light transport*, *photon mapping* and *instant radiosity* [VG97; Jen96; Kel97]. Among these methods, instant radiosity stands out as an approach that easily maps to standard rasterization APIs, since all indirect illumination is approximated by VPLs. To compute their distribution, initial light sources emit multiple VPLs, whose path is traced throughout the scene. Each intersection with geometry may lead to another bounce or the termination of the path and final placement of the VPL. Based on this idea, many-light rendering has emerged as a new class of algorithms for image synthesis [Dac+14; SIP07; Rit+08; LY19]. While using a higher number of VPLs yields superior results in most many-light methods, the implied effort for shading severely inhibits scalability.

## 2 Background and Related Work

To cap the number of VPLs that participate in shading, Walter et al. [Wal+05] presented *lightcuts*. They generate a static binary light tree over all VPLs, where each inner node represents the lights in its leaf. These *representative lights* are obtained by recursively merging nodes with similar attributes. While the position of each representative is chosen randomly from one of its child nodes, its intensity equates to their sum to satisfy the preservation of radiant flux in arbitrary tree cuts. For a portion of the output image domain (e.g., one pixel), they traverse the light tree and identify a lightcut that bounds the error caused by shading with the corresponding representatives. Visibility determination plays a key role, because fully occluded light sources and subtrees can be trivially rejected. Multidimensional lightcuts [Wal+06] extend the core algorithm to raise its efficiency in higher-order domains accounting for visual effects like depth-of-field, motion blur and participating media.

Due to the inherent sample correlation between cuts from the static light tree, visible indicators of undersized lightcuts usually manifest as discretized penumbrae and isolated “blobs” of light, making them difficult to resolve with automated filtering. Recently proposed stochastic lightcuts by Yuksel [Yuk19] trade these structured artifacts for more Monte Carlo noise. The selection of inner nodes’ light positions is performed probabilistically via importance sampling of their respective leafs in each individual lightcut, thus rejecting the idea of precomputed representatives.

As lightcuts were originally devised as an offline method on the CPU, Hašan, Pellacini, and Bala proposed a GPU-friendly, matrix-based method for scalable many-light rendering [HPB07]. Formulating the problem as an  $M \times N$  matrix with  $M$  pixels and  $N$  light sources, they sample and select columns with high overall contribution and compute their weighted sum to obtain high-quality images at interactive rates. Several optimizations have been proposed that extend the original idea and further improve its performance [OP11; Huo+15]. However, while faster than lightcuts, reported runtimes for rendering single frames with these techniques remain in the order of seconds to minutes.

Stochastic methods for image synthesis can often be enhanced by the application of reconstruction techniques on noisy or undersampled input. In

many cases, a single sample per pixel suffices to produce high-fidelity output with state-of-the-art image reconstruction techniques [LWC12; Dam+10]. In this context, animations pose a particular challenge since reconstruction must take both spatial and temporal stability into account to preserve frame-to-frame coherence [Sch+17; SPD18]. Recent work focusing on machine learning has shown that spatio-temporal denoising can largely be delegated to the run-time inference of pre-trained neural networks [Cha+17].





## 3 Stochastic Substitute Trees

Our complete pipeline is outlined in Fig. 3.1: We first distribute VPLs according to IR and build our SST over the resulting VPLs. We then generate G-buffers using rasterization and for all fragments sample the illumination from the SST, yielding a noisy light image. With the help of G-buffers, we use a CNN for denoising and temporal filtering to ensure consistent shading over time. In the following, we give a detailed description of the substitute tree concept and its parallel construction. In Section 4, we describe the process of drawing samples from the SST and provide additional, implementation-specific details in Section 5.

### 3.1 The Substitute Tree

Our substitute light trees are motivated by the original methods for lightcuts and, specifically, stochastic lightcuts [Wal+05; Yuk19]. The inner nodes of the light trees used in these approaches copy their geometric properties from one of the leaf nodes, *i.e.*, one of the original VPLs. This policy necessarily leads to noticeable lighting patterns when the number of VPLs is low. In order to support dynamic light sources and animated scenes in real-time, we aim to keep the number of VPLs low to permit updating the light tree in every frame, while minimizing visual artifacts.

To better approximate the sampling of the—in general smooth—indirect illumination, we do not use leaf VPL locations for the inner nodes of our light tree. Instead, we pretend as if the leaves were distributed more densely than they actually are by first merging VPLs into clusters and then approximating those clusters with fitting *substitute distributions* (see Fig. 3.2). Although we could theoretically use any probability distribution to describe the clusters,

### 3 Stochastic Substitute Trees

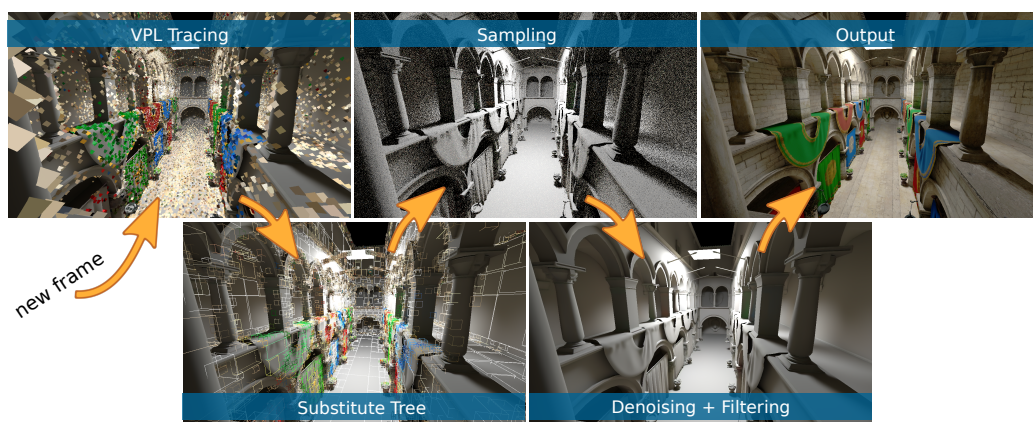



Figure 3.1: Global illumination with stochastic substitute trees on a GeForce RTX 2080Ti in  $1280 \times 720$ : For every frame, we distribute VPLs (0.4 ms), build a new substitute tree over the VPLs (0.9 ms), stochastically sample the substitute tree (3.4 ms), apply CNN denoising (17.4 ms) and temporal filtering (0.9 ms). Substitute trees use random distributions as inner nodes which allow efficient sampling and approximate the light distributions as if larger numbers of VPLs were used.


we aim to use a compact and efficient representation for sampling. We intuitively choose a normal distribution to describe the individual clusters, as the central limit theorem suggests that the addition of independent random variables will take the shape of a normal distribution. Note that the central limit theorem only partially holds as the light distributions are in general not independent.

Similar to previous approaches, we construct the full light tree bottom-up and merge substitute clusters as we go. For traditional lightcuts, ideal merge candidates should be spatially close and have similar normals to minimize the view-dependent deviations from full shading when using representatives. For our approach, the situation is more complicated since we must ensure that estimated distributions align with the geometry of the scene. For instance, a naïve substitute distribution may include portions of the scene that are occluded or “outside” of the scene. As describing distributions that follow non-trivial surfaces would be costly, we try to limit clusters to surfaces that can be described and sampled straightforwardly. The simplest surface, a plane, can be described by a location and normal. Incidentally, this corresponds to the required information for shading with a VPL: its

### 3.1 The Substitute Tree



(a) Traced VPLs



(b) Substitute Distributions

Figure 3.2: (a) 150 k VPLs distributed in the Pink Room with their RGB intensity. (b) While previous work uses discrete locations from the input VPLs for inner tree nodes, our stochastic substitute trees store continuous distributions of VPLs, defined by a substitute position and variance: red and green lines show the standard deviations on the surface; blue lines along the normal. Notice how large planar regions are captured by single distributions and variance along the normal is virtually undetectable.

location and the surface normal. Following our previous considerations, we generate normally distributed samples on the plane defined by the cluster position and normal. This strategy works well for nearly planar surfaces, but surfaces with high curvature or corners are difficult to describe. However, the memory requirement is minimal, in fact, identical to Lightcuts. The alternative, a multivariate normal distribution would also work but has greater memory and computational requirements.

## 3.2 Distribution Merging and Errors

To obtain a suitable substitute distribution for each inner node, we compute the intensity-weighted cluster position  $\mathbf{p}$  and the positional variance  $\sigma^2$ . To have  $\sigma_x^2$  and  $\sigma_y^2$  represent the variance of lights on the plane surface, we transform their positions into a local plane coordinate system where the z-axis matches the cluster plane normal. Merged position and variance are computed using

$$\mathbf{p} = \frac{c_1 \cdot \mathbf{p}_1 + c_2 \cdot \mathbf{p}_2}{c_1 + c_2} \quad \text{and} \quad (3.1)$$

$$\sigma^2 = \frac{c_1 \cdot (\sigma_1^2 + \Delta'_{p1} \odot \Delta'_{p1}) + c_2 \cdot (\sigma_2^2 + \Delta'_{p2} \odot \Delta'_{p2})}{c_1 + c_2}, \quad (3.2)$$

where  $c_1, c_2$  are the child contributions,  $\mathbf{p}_1, \mathbf{p}_2$  are the child cluster positions,  $\sigma_1^2, \sigma_2^2$  are the child cluster variances, and  $\Delta'_{p1}, \Delta'_{p2}$  are the difference vectors from  $\mathbf{p}$  to  $\mathbf{p}_1$  and  $\mathbf{p}_2$  in transformed cluster space. The weights  $c_1, c_2$  correspond to the relative luminance of the cluster's RGB light intensity. While we could compute a spread (and thus variance) alongside the VPL distribution and use this spread as initial variance, we have found that starting with simple VPLs works equally well in practice and avoids complex spread computations. An example for a substitute tree is given in Fig. 3.2b.

Similar to regular Lightcuts, the merged node's intensity is simply the sum of the child intensities and its extent is encoded by an axis-aligned bounding box that encloses all light sources in its subhierarchies. The merged normal is linearly interpolated based on the child contributions  $c_1, c_2$ . Hence, clusters with high weights (and thus high intensities) have a stronger influence on the outcome.

Note that strongly diverging normals will alter the cluster coordinate system and thus partially invalidate the local variance values during merging. To avoid corresponding artifacts, we therefore compute and store additional information on each cluster's suitability for sampling. During sampling, we then skip any distribution that does not fulfill the corresponding criteria. For evaluating the suitability of a substitute node for sampling, we use three indicators:

- The normal similarity metric  $\nu = \nu_1 \cdot \nu_2 \cdot \max(\langle \mathbf{n}_1, \mathbf{n}_2 \rangle, 0)$

- The variance  $\sigma_z$  along the plane normal
- The diagonal size  $\tau$  of the axis-aligned bounding box

For the normal similarity metric  $\nu$ ,  $\mathbf{n}_1$  and  $\mathbf{n}_2$  as well as  $\nu_1$  and  $\nu_2$  capture normals and the normal similarities of the child clusters. The normal similarity metric serves as an indicator for both the similarity between the child cluster distribution orientations and the curvature of the surface. Furthermore,  $\nu$  is guaranteed to decrease when going up the tree. The diagonal size  $\tau$  of the bounding box limits the size of the plane described by our substitute distribution. Since we do not use color similarity, this indicator helps to avoid merging different color patches. For leaf nodes the normal similarity is initialized to one and the variance to zero. Given a set of suitable parameter thresholds, we can combine these indicators to form a joint condition for ensuring that merged substitute clusters are of sufficient quality for sampling:

$$\nu \geq \nu_{\min} \wedge \sigma_z \leq \sigma_{\max} \wedge \tau \leq \tau_{\max}. \quad (3.3)$$

### 3.3 Parallel Construction

Since we aim to rebuild the light tree in every frame, minimizing its construction time is essential. Traditional lightcuts follow an  $\mathcal{O}(n^2)$  approach to find pairs of lights/clusters that minimize the introduced error according to their metric. As this approach is incompatible with our real-time constraints, we follow the algorithm by Karras [Kar12] for building bounding volume hierarchies.

To start the hierarchy construction, we sort VPLs according to a spatial indexing using Morton codes. Additionally, we incorporate bits extracted from the normals, since we seek VPLs that point in similar directions to obtain candidates for suitable substitute distributions. Using polar coordinates to describe normal directions leads to non-uniform quantization, as bits cover differently-sized areas on the sphere. Thus, we first identify the octant the normal points to and then subdivide the octant as a triangle to yield approximately equally-sized areas [TR98], as shown in Fig. 3.3. Each octant is represented by three bits, followed six bits to describe the

### 3 Stochastic Substitute Trees

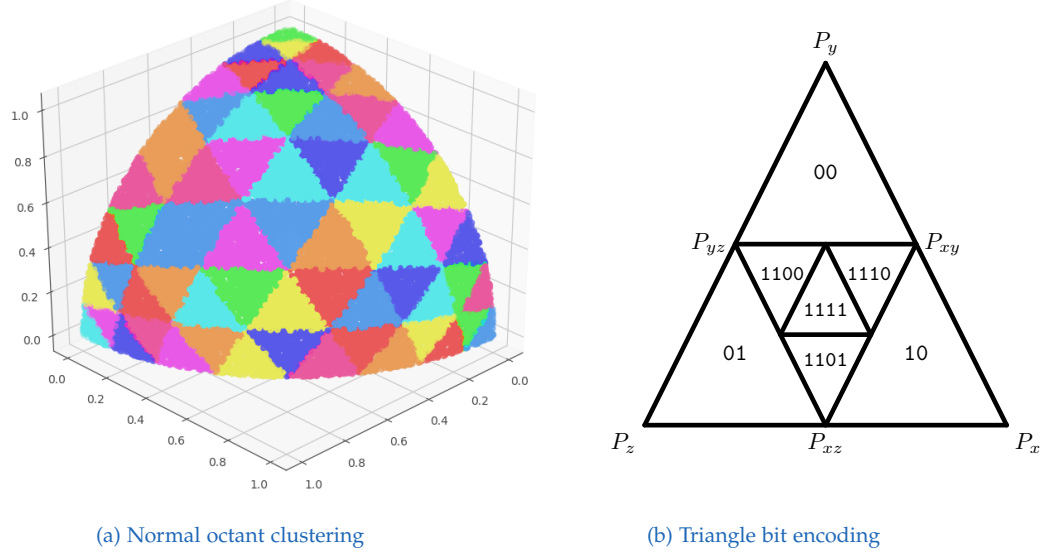


Figure 3.3: (a) Normal clustering in an octant of the unit sphere using 3 refinement levels; (b) the corresponding bit values for two refinements.

subdivision level of the triangle. Our final code  $C$  for each VPL is composed of the Morton code  $M$ , normal code  $N$  and a unique id:

$$C = M \oplus N \oplus id, \quad (3.4)$$

where  $\oplus$  is the bit concatenation operator. We thus sort primarily according to position, since VPLs lying on the same surface will be close to each other. Secondary sorting occurs according to normals, to ensure that VPLs with similar normals are prioritized merge targets. To circumvent any duplicate codes, a unique ID is concatenated at the end. In the final 64-bit code,  $M$  uses 30 bits and  $N$  uses 9 bits. The remaining bits are used for the unique id.

Following the initial sort according to (3.4), the tree is constructed bottom-up. For merging the inner node substitute distributions, one thread per leaf node is started. At each level, only one thread may continue to the next-higher parent. After computing the distributions and corresponding quality metrics, the node information is stored as an array of structs. For each node, we store:

### 3.3 Parallel Construction

- distribution variance, ID and child IDs, (3 floats + 3 ints)
- the world space position and normal, ( $2 \times 3$  floats)
- the bounding box, ( $2 \times 3$  floats)
- light color and intensity, (4 floats)
- a flag indicating if the distribution is suitable for sampling,

yielding a total structure size of 96 bytes (padded). Thus, the full node information can be written with six vector store instructions.





## 4 Stochastic Light Sampling

We sample our stochastic substitute tree by traversing it until we arrive either at a leaf VPL or a suitable inner node, in which case we draw the sample from its substitute cluster distribution. For drawing multiple samples, we can simply run the traversal multiple times from the root. However, since the evaluation of a sample requires expensive visibility determination, we aim to maximize the quality gain using only one sample per pixel.

### 4.1 Tree traversal

We define a cluster's illumination  $L_C$  for a shading point  $x$  from direction  $\omega$  with the same terms as in the original lightcuts paper:  $\omega$

$$L_C(x, \omega) \approx M_j(x, \omega) \cdot G_j(x) \cdot V_j(x) \cdot I_j, \quad (4.1)$$

where for a tree node  $j$ ,  $M_j$  is the material term,  $G_j$  is the geometric term,  $V_j(x)$  represents visibility of  $x$  from the light source and  $I_j$  is the sum of VPL intensities in the subtree. In order to stochastically sample clusters based on their expected contribution to the shading of  $x$ , we start traversal from the root and continually compute weights  $w_1, w_2$  for the current node's children. Similar to stochastic lightcuts, the probability of choosing the  $i$ th child is given by  $\frac{w_i}{w_1+w_2}$ . If we enter a *dead branch* (i.e.,  $w_1 + w_2 = 0$ ), we stop the traversal. While this wastes a light sample, backtracking within the shader is not feasible for our real-time performance requirements. We compute  $w = \mathcal{O}(L_C)$  for each child by plugging in adequate upper bounds for those terms in (4.1) that are too expensive to evaluate. To map the RGB value of  $I_j$  to  $\mathbb{R}^1$ , we use its relative luminance in practice.  $V_j$  can be trivially set to 1, thus assuming no occlusions.

## 4 Stochastic Light Sampling




Figure 4.1: Comparing the accuracy of color bleeding in the final image. While SST (b) faithfully captures the soft reflection of the red fabric on the pillar as it is seen in the ground truth (a), it is almost entirely missing with SLC (c). The bottom right half of (b,c) shows the pixel difference  $\times 10$ .

For the material term  $M_j$ , we use the upper bound of the cosine between  $\omega$  and the surface normal at  $x$ , multiplied by the BRDF. The upper bound of the cosine is computed w.r.t. the bounding box [Wal+05]. The diffuse BRDF is independent of the view or light direction, thus constant. For the specular part of the BRDF, we use the following heuristic: We compute the perfect reflection  $\mathbf{R}$  of the view vector along the surface normal at  $x$ , *i.e.*, a mirror-like reflection. We then test if the ray  $x + t \cdot \mathbf{R}$  intersects with the cluster’s bounding box. If so, we use  $\mathbf{R}$  for the specular BRDF computation. Otherwise, the BRDF is evaluated with a direction vector  $\mathbf{R}'$  that actually does intersect the bounding box and minimizes the difference to  $\mathbf{R}$ .

Since we focus on omni-directional lights,  $G_j$  reduces to inverse quadratic attenuation. We note that, in order to guard against singularities in the geometric term, stochastic lightcuts set  $G_j = 1$  for omni-directional lights if  $x$  lies inside the cluster’s bounding box, *i.e.*, ignoring the distance from the sample to the cluster. However, we found that this leads to missing color bleeding, as shown in Fig. 4.1. Noticeable color bleeding at  $x$  is often due to VPLs that are close to  $x$  and thus—especially for lower levels of

## 4.2 Sampling the Substitute Distribution




Figure 4.2: (left) The color-coded tree traversal depth ( ■ / ■ for low/high) is on average 20 for SLC, whereas SST only requires 15 steps on average for a 100 k VPL tree. (right) SST are both faster and achieve better image quality compared to the ground truth (difference  $\times 10$ ).

the tree—likely to be enclosed by the same bounding box as  $x$ , leading to rejection of the distance term. This makes tree traversal more likely to end up in different subtrees, which do not contribute to the color bleeding. In contrast, we avoid singularities by simply bounding the distance from below against an  $\epsilon$  term to guard against division by zero.

## 4.2 Sampling the Substitute Distribution

In contrast to previous many-light methods, which usually focus on diffuse shading only, we also consider non-Lambertian BRDFs. Before traversing the tree, we decide whether to sample the diffuse or specular portion of the BRDF, based on the magnitude of the respective material terms. For diffuse samples, we stop the traversal at the first usable substitute distribution, *i.e.*, the first node that has been marked suitable for sampling. A suitable substitute distribution is often found early on, reducing the sampling depth significantly without any visible difference to the final illumination (see Fig. 4.2). For drawing a light position sample  $p_S$  from the distribution, we generate three normally distributed random variables, according to the cluster variance  $\sigma^2$  in the local coordinate frame. After transformation to world coordinates, we clamp  $p_S$  by the node’s bounding box and evaluate the incoming radiance according to Eq. (4.1). This implies accurate visibility

## 4 Stochastic Light Sampling

testing to determine the binary value for  $V_j$  at shading point  $x$ , which we resolve through ray tracing.

When sampling the specular portion of the BRDF, we ignore the substitutes and only consider leaf VPLs as viable candidates, since highly specular reflections would be overly blurred by picking substitute samples. While this leads to performance penalties for specular surfaces, it is vital for ensuring high quality.

## 5 Implementation

We implemented our approach on top of the NVIDIA Falcor framework [Ben+19] in DirectX 12 and CUDA. Raytracing is done via hardware accelerated RTX. For diffuse lighting, we use the Lambertian model. For the specular BRDF, we use Cook-Torrance reflectance [CT82].

*VPL Generation.* For VPL generation, we follow the core approach described by instant radiosity [Kel97]. We start by probabilistic sampling of the original light sources according to their size and intensity. The number of VPL paths to trace for each light depends on its intensity, as well as the configured maximum number of bounces. To trace these rays, we use hardware supported DX12 raytracing. At every hit, we generate a new VPL and append it to a global buffer. We use cosine-weighted sampling for the next bounce direction and early path termination via Russian Roulette.

*Substitute Tree Generation.* Substitute trees are built in parallel on the GPU with CUDA. Internal nodes are stored in the same buffer as the previously generated VPLs. The tree is constructed by a single kernel, decreasing the number of active threads while progressing up through the tree. Since the number of active nodes usually quickly falls below the core count of modern GPUs, the performance loss due to keeping inactive threads around is minimal.

*Sampling.* Substitute tree sampling is directly done by a ray-tracing shader, since we use shadow rays for visibility. We first fill a GBuffer with first hit shading information [Dee+88]. For each pixel, we traverse the tree to draw a VPL sample and cast a shadow ray from the VPL location  $p_S$  to the pixel's stored position in world space in the GBuffer. We found that tracing in this direction rather than from  $x$  to  $p_S$  yields better performance, which we ascribe to increased coherence achieved in supplying multiple pixels with similar VPL samples. To account for direct illumination, we randomly pick

## 5 Implementation

and sample one of the scene’s light sources. Thus, we trace two shadow rays per pixel—one for tree sampling and one for direct illumination.

*Denoising and Temporal Stability.* Our approach produces Monte Carlo noise and thus, in theory, any Monte Carlo denoiser is applicable. In practice, we use a pruned recurrent autoencoder (RDAE) with the same training procedure and loss function as stated in the original paper [Cha+17]. To further improve temporal stability, we apply the temporal accumulation step of the adaptive spatio-temporal variance-guided filter [Sch+17; SPD18], using the pixel-wise intensity gradient from the RDAE output as temporal accumulation factor to control blending. This mitigates artifacts like ghosting and temporal lags, especially for animated scenes.

To save time on inference, we reduce the number of features of RDAE, while still maintaining high quality. For downsampling, we use 24, 40, 54, 72, 96, and 96 features. For upsampling, we use (96, 96), (72, 72), (54, 54), (40, 40), (64, 32). As activation function we use ReLu, for upsampling we use the deconvolution operator. The network was modeled and trained with Tensorflow and TensorRT as inference engine. The training data was generated with the Sponza and San Miguel scene, with 500 captured frames per scene. Inference is executed on NVIDIA tensor cores in fp16 mode.

## 6 Evaluation

We compare our stochastic substitute trees (SST) in detail to regular path tracing (PT) with varying number of bounces, stochastic lightcuts (SLC), and show the main differences to the recently proposed real-time method for Lighting Grid Hierarchies [LY19] (LGH). All timings are generated on a GeForce RTX 2080Ti. We use  $\nu_{\min} = 0.5$  and  $\sigma_{\max} = 0.1$ ; and  $\tau_{\max} = \infty$ , as we found no benefits of limiting the cluster size for the relatively small test scenes. The reference path tracer is based on the NVIDIA Falcor framework. The first hit is determined by a rasterized G-Buffer, additional bounces are computed with RTX. *Next event estimation* is applied to improve convergence. For quality comparisons, we use path tracing with 4000 samples per pixel (spp) and seven bounces as ground truth.

We evaluate on four scenes with an image resolution of  $1280 \times 720$  shown in Fig. 6.1. The Pink Room and Living Room are simple in terms of geometry, but offer interesting lighting situations and highly glossy surfaces. Sponza is slightly more complex in geometry and shows high amounts of color bleeding and shadowing through the corridors. Finally, San Miguel offers high geometric load and enforces diverse light paths. Fig. 6.1 shows that SST produces lighting close to the ground truth and accurately reconstructs the lighting under noisy conditions and complex setups. Some difficult-to-reach areas, like the table in the living room scene or the shelves in the pink room, appear slightly darker, which is typical for VPL approaches.

For runtime evaluation, we separately look at SST construction, which we can run in every frame or when lighting changes; SST sampling, which rasterizes a G-buffer, samples the tree and performs direct lighting using RTX; denoising; and temporal filtering. For temporally stable lighting, we require  $\sim 100$  k VPLs. Their distribution takes 0.27 ms to 0.85 ms, with an additional 0.9 ms for SST building. Sampling takes 2.5 ms to 7.7 ms, denoising 17.4 ms




Figure 6.1: Sampling result of SST with 1 spp for rook VPLs, denoised and ground truth comparison. While there are some areas (e.g. the wall in the Pink Room) where our approach misses some light, SST overall produces high quality lighting.



and temporal filtering 0.9 ms. Including VPL distribution and tree building, an entire frame takes 21.97 ms to 27.75 ms. While we achieve 30fps overall, denoising is clearly the limiting factor, making up 63 % to 80 % of frame time. With GPU hardware set to further advance CNN performance, we expect this issue to lessen in the future.

Table 6.1 shows the sample timings and final image quality for all tested approaches. SST is always faster and always achieves superior quality to SLC. We attribute our quality advantage to two factors: First, using substitute distributions approximates global illumination as if it were generated with more VPLs and thus can achieve image qualities for which other approaches require significantly higher VPL counts. Second, our node selection during tree traversal generalizes attenuation also to close-by VPLs, which have a higher contribution to localized lighting effects (see Fig. 4.1).

## 6.1 Comparison

Since the ground truth is generated with path tracing, PT variants may be slightly advantaged w.r.t. the quality comparison. Nevertheless, Table 6.1 shows that while PT with a single bounce is efficient, the image quality is clearly inferior to SST (see also Fig. 6.2). In all scenes, there are areas that require more than a single bounce to be properly lit. When the number of bounces is increased, PT catches up in terms of quality, but also requires more time than ours, *i.e.*, PT requires 2 to 3 bounces to achieve our quality while being  $1.5\times$  slower. San Miguel is an exception, where all approaches struggle to achieve good quality, showing the highest error among all scenes. However, substitute trees are faster than even single-bounce PT, as ray tracing is particularly expensive due to the geometric load.

To quantify the potential gains of the different approaches, we list the SSIM gain per invested ms when switching to either technique from the lowest-quality approach, *e.g.*, switching from  $PT_1'$  to SST for Pink Room, achieves an  $\sim 0.052$  SSIM increase for the invested 1 ms. SST always yields the best gains for pure diffuse lighting. When using a non-Lambertian BRDF, scenes with highly glossy materials, like Pink Room and Living Room, work slightly better with path tracing. Sponza, featuring only moderately glossy

	Sampling time (ms)			RMSE denoised			SSIM denoised			Gain (% SSIM/ms)										
	SST	SLC	PT <sub>g</sub>	PT <sub>1'</sub>	PT <sub>2'</sub>	PT <sub>g'</sub>	SST	SLC	PT <sub>g'</sub>	PT <sub>1'</sub>	PT <sub>2'</sub>	PT <sub>g'</sub>	SST	SLC	PT <sub>1'</sub>	PT <sub>2'</sub>	PT <sub>g'</sub>			
Pink Room	2.65	3.59	1.73	3.20	4.71	0.36	0.48	0.133	0.81	0.53	0.60	0.951	0.912	0.956	0.968	5.22	2.10	0.00	2.99	1.88
Living Room	2.88	4.34	1.78	3.33	4.76	0.26	0.33	0.070	0.39	0.24	0.60	0.956	0.928	0.965	0.972	2.91	1.09	0.00	2.39	1.48
Sponza	3.36	4.54	2.75	5.45	8.10	0.17	0.17	0.029	0.15	0.12	0.71	0.967	0.877	0.967	0.975	15.41	5.03	0.00	3.33	1.83
San Miguel	5.77	6.83	8.86	17.81	27.18	0.48	0.49	0.046	0.45	0.45	0.872	0.862	0.873	0.880	0.881	1.89	0.00	0.54	0.16	0.09
Pink Room	4.24	5.08	1.88	3.62	5.33	0.48	0.59	0.147	0.92	0.62	0.947	0.940	0.900	0.943	0.958	1.99	1.25	0.00	2.47	1.68
Living Room	4.11	5.14	1.91	3.77	5.27	0.40	0.47	0.087	0.51	0.32	0.946	0.940	0.910	0.952	0.966	1.64	0.93	0.00	2.26	1.67
Sponza	3.99	5.31	2.95	5.82	8.67	0.17	0.17	0.029	0.15	0.12	0.970	0.967	0.876	0.967	0.975	9.04	3.86	0.00	3.17	1.73
San Miguel	7.34	7.84	8.76	17.83	27.45	0.72	0.74	0.065	0.59	0.58	0.758	0.752	0.772	0.793	0.795	2.40	0.00	2.17	0.41	0.22

Table 6.1: Comparison between SST and SLC (took VPLs) and path tracing with increasing number of bounces. While 1 bounce PT is faster than SST, its quality is lower, as it misses large amounts of indirect lighting. With increasing bounces, PT achieves comparable quality, but is considerably slower. To quantify the potential performance gains, we show the relative increase SSIM per invested ms compared to the worst performing approach (higher is better, best in bold).



Figure 6.2: Path tracing with a single bounce (left) misses significant portions of indirect illumination. SST (right) distributes VPLs throughout multiple bounces (in a fraction of the trace time) and thus faithfully recovers the illumination.

materials, again works best with SST. In San Miguel, little gains can be achieved overall with either method.

Unfortunately we were unable to include LGH in this direct comparison as the LGH framework is not compatible with Falcor to produce identical scene setups. Nevertheless, we compare a view of Sponza directly between LGH and SST, as shown in Fig. 6.3: It is lit with indirect illumination generated from a single point light, placed at the center of the scene. The ground truth is generated by the LGH framework using instant radiosity. It is evident that LGH significantly overestimates lighting in some places, while being too dark in others. We attribute these issues to LGH’s shadow heuristic, which is insufficient to distinguish between the brighter (but occluded) VPLs in the center of the scene and the less bright (but visible) VPLs in the corridors. Increasing the shadow ray count to 2 per pixel (the maximum supported by the framework) did not improve quality, but prolonged the frame time from 26.7 ms to 37.6 ms (from hierarchy generation to final image). For the entire pipeline, SST requires 24.57 ms and creates clearly superior results compared to LGH.

## 6 Evaluation




Figure 6.3: While LGH requires similar times as SST, it is further from the ground truth indirect illumination. LGH’s shadowing heuristics overestimate the general illumination and predict high shadows around corners (red boxes).

### 6.2 Scalability

Since we regularly (potentially in every frame) generate new VPLs and the substitute tree, performance of these steps is essential. Table 6.2 shows the timings for VPL generation and tree building with different scenes. The runtime for VPL generation is governed by the scene’s complexity, while the duration of tree building depends only on the number of VPLs. About 100 k VPLs are needed for temporally stable lighting in our tests, resulting in  $\sim 1.1$  ms to 1.5 ms in total.

Table 6.3 shows that the sample time also increases with tree depth. Sampling from the SST is more efficient as soon as the tree is not too shallow. For very small VPL counts (10k)—which imply temporally unstable lighting conditions—SLC is slightly faster. For 100k VPLs, SST is already  $1.35\times$  faster. Our performance edge increases with the tree size, indicating a better scalability of substitute trees. When sampling specular reflections, the perfor-

## 6.2 Scalability

# VPLs	VPL generation				Tree building			
	10k	100k	500k	1M	10k	100k	500k	1M
Pink Room	0.13	0.27	0.74	1.47	0.58	0.89	2.44	4.51
Living Room	0.13	0.30	0.83	1.54	0.54	0.89	2.38	4.43
Sponza	0.14	0.40	1.32	2.38	0.59	0.88	2.39	4.35
San Miguel	0.19	0.85	3.44	7.05	0.56	0.91	2.35	4.38

Table 6.2: Timings for VPL generation and substitute tree building with varying VPL counts. VPL generation depends on the scene geometry due to ray-tracing. Tree building only depends on the VPL count. All times in ms.

		10k		100k		500k		1M	
		SST	SLC	SST	SLC	SST	SLC	SST	SLC
diffuse	Pink Room	2.22	<b>2.10</b>	<b>2.65</b>	3.59	<b>2.88</b>	5.45	<b>3.10</b>	6.79
	Living Room	2.27	<b>2.17</b>	<b>2.88</b>	4.34	<b>3.36</b>	6.63	<b>3.73</b>	8.04
	Sponza	2.77	<b>2.74</b>	<b>3.36</b>	4.54	<b>3.78</b>	6.12	<b>3.92</b>	7.00
	San Miguel	<b>4.81</b>	5.19	<b>5.77</b>	6.83	<b>6.64</b>	8.51	<b>7.30</b>	9.39
specular +	Pink Room	3.90	<b>3.62</b>	<b>4.61</b>	5.08	<b>5.02</b>	7.08	<b>5.18</b>	9.14
	Living Room	3.68	<b>3.27</b>	<b>4.43</b>	5.14	<b>5.03</b>	7.01	<b>5.24</b>	8.86
	Sponza	3.36	<b>3.26</b>	<b>4.10</b>	5.31	<b>4.58</b>	6.99	<b>4.73</b>	8.02
	San Miguel	6.45	<b>6.29</b>	<b>7.71</b>	7.84	<b>8.92</b>	9.85	<b>10.04</b>	11.80

Table 6.3: Sampling timings in ms with 1spp. For shallow trees (which are temporally not stable) SLC is slightly faster than sampling from SST. However, as the tree size increases, our substitute sampling clearly has the performance edge.

mance difference is less pronounced, as specular sampling is more complex and we traverse the tree further. Nevertheless, SST is still faster for all tree sizes above 10k VPLs. Our performance edge comes from two factors. First, we stop our traversal early, as outlined in Fig. 4.2. Second, SLC’s random sampling of VPL locations leads to scattered memory accesses, while SST sampling leads to coherent fetches of only the distribution information and thus requires less memory bandwidth.

## 6 Evaluation




Figure 6.4: (a) An area light source, pointed at a corner, is sampled well with SST (left). Path tracing (right) has difficulties connecting the rays to the light source, which leads to undersampling. (b) While SST (left) supports specular BRDFs, highly glossy surfaces are challenging to resolve with 1spp, as the reflection directions are usually not sampled and thus more noise remains compared to path tracing (right).

### 6.3 Discussion and Limitations

So far, we only reported our results on image quality for individual still frames. However, our experiments have shown that denoising the result of our sampling stage with temporal filtering produces temporally stable global illumination in animated scenes, where we rebuild the SST each frame with sufficiently many VPLs. Thus, our approach supports dynamic light sources and scenes while maintaining good frame-to-frame coherence. Due to the difficulties of establishing temporal stability via conventional illustrations and for lack of a universally accepted quality metric, we kindly refer the reader to our supplemental video for an adequate demonstration.

Compared to path tracing, one particular advantage of SST is that it can handle difficult-to-reach light sources due to its VPL distribution in the style of instant radiosity, as shown in Fig. 6.4a.

Although SST does support specular BRDFs, mirrors and glossy reflections pose a challenge, as sampling the tree will in general not sample the perfect reflection direction and thus degrade sharpness (see Fig. 6.4b). As a remedy, one can trace an additional ray for samples that hit a mirror and sample the tree from the next hit.

### 6.3 Discussion and Limitations

Working with VPLs, using heuristically driven importance sampling, and estimating VPL distributions leads to a biased approach. Additionally, as is common with VPL approaches, we avoid singularities in the geometric term by clamping, which results in slightly darkened corners. However, focusing on the creation of convincing real-time rendering content, being unbiased is not our main goal.





## 7 Conclusion

We have presented a new, scalable many-light rendering approach for simulating global illumination in real-time in animated scenes. Our stochastic substitute trees approximate the actual global illumination by converting clusters of input VPLs into suitable probability distributions. While we can benefit from the same advantages as stochastic lightcuts, sampling with our tree is more efficient since traversal can be stopped early. Furthermore, our substitute distributions generate higher-quality lighting than approaches that directly sample from the underlying VPLs. Our sampling method compares favorably to stochastic lightcuts and the recently proposed Lighting Grid Hierarchy in terms of quality and runtime. We can usually capture indirect illumination more accurately with a single bounce than path tracing and report higher gains for image quality metrics. By combining our sampling method with CNN denoising and temporal filtering, we can generate temporally stable approximations of global illumination in real-time, in fully dynamic scenes.

Our approach supports specular highlights and produces agreeable results for semi-glossy surfaces. However, highly glossy surfaces and sharp, mirror-like reflections may still produce noise and are better handled with path tracing. Looking forward, we believe that tracing additional rays for those hits may combine the best of both worlds, leading to hybrid VPL sampling and path tracing.



# Bibliography

- [AL09] Timo Aila and Samuli Laine. “Understanding the Efficiency of Ray Traversal on GPUs.” In: *Proceedings of the Conference on High Performance Graphics 2009*. HPG ’09. New Orleans, Louisiana: ACM, 2009, pp. 145–149. ISBN: 978-1-60558-603-8. DOI: [10.1145/1572769.1572792](https://doi.org/10.1145/1572769.1572792). URL: <http://doi.acm.org/10.1145/1572769.1572792> (cit. on p. 3).
- [Ben+19] Nir Benty et al. *The Falcor Rendering Framework*. <https://github.com/NVIDIAGameWorks/Falcor>. Oct. 2019. URL: <https://github.com/NVIDIAGameWorks/Falcor> (cit. on p. 19).
- [Cha+17] Chakravarty R. Alla Chaitanya et al. “Interactive Reconstruction of Monte Carlo Image Sequences Using a Recurrent Denoising Autoencoder.” In: *ACM Trans. Graph.* 36.4 (July 2017), 98:1–98:12. ISSN: 0730-0301. DOI: [10.1145/3072959.3073601](https://doi.org/10.1145/3072959.3073601). URL: <http://doi.acm.org/10.1145/3072959.3073601> (cit. on pp. 5, 20).
- [CT82] R. L. Cook and K. E. Torrance. “A Reflectance Model for Computer Graphics.” In: *ACM Trans. Graph.* 1.1 (Jan. 1982), pp. 7–24. ISSN: 0730-0301. DOI: [10.1145/357290.357293](https://doi.org/10.1145/357290.357293). URL: <http://doi.acm.org/10.1145/357290.357293> (cit. on p. 19).
- [Dac+14] Carsten Dachsbacher et al. “Scalable Realistic Rendering with Many-Light Methods.” In: *Comput. Graph. Forum* 33.1 (Feb. 2014), pp. 88–104. ISSN: 0167-7055. DOI: [10.1111/cgf.12256](https://doi.org/10.1111/cgf.12256). URL: <http://dx.doi.org/10.1111/cgf.12256> (cit. on p. 3).

## Bibliography

- [Dam+10] Holger Dammertz et al. “Edge-avoiding À-Trous Wavelet Transform for Fast Global Illumination Filtering.” In: *Proceedings of the Conference on High Performance Graphics*. HPG '10. Saarbrücken, Germany: Eurographics Association, 2010, pp. 67–75. URL: <http://dl.acm.org/citation.cfm?id=1921479.1921491> (cit. on p. 5).
- [Dee+88] Michael Deering et al. “The Triangle Processor and Normal Vector Shader: A VLSI System for High Performance Graphics.” In: *SIGGRAPH Comput. Graph.* 22.4 (June 1988), pp. 21–30. ISSN: 0097-8930. DOI: [10.1145/378456.378468](https://doi.org/10.1145/378456.378468). URL: <http://doi.acm.org/10.1145/378456.378468> (cit. on p. 19).
- [HA19] Eric Haines and Tomas Akenine-Moller. *Ray Tracing Gems: High-Quality and Real-Time Rendering with DXR and Other APIs*. Berkely, CA, USA: Apress, 2019. ISBN: 1484244265, 9781484244265 (cit. on p. 3).
- [HPB07] Miloš Hašan, Fabio Pellacini, and Kavita Bala. “Matrix Row-column Sampling for the Many-light Problem.” In: *ACM Trans. Graph.* 26.3 (July 2007). ISSN: 0730-0301. DOI: [10.1145/1276377.1276410](https://doi.org/10.1145/1276377.1276410). URL: <http://doi.acm.org/10.1145/1276377.1276410> (cit. on p. 4).
- [Huo+15] Yuchi Huo et al. “A Matrix Sampling-and-recovery Approach for Many-lights Rendering.” In: *ACM Trans. Graph.* 34.6 (Oct. 2015), 210:1–210:12. ISSN: 0730-0301. DOI: [10.1145/2816795.2818120](https://doi.org/10.1145/2816795.2818120). URL: <http://doi.acm.org/10.1145/2816795.2818120> (cit. on p. 4).
- [Jen96] Henrik Wann Jensen. “Global Illumination Using Photon Maps.” In: *Proceedings of the Eurographics Workshop on Rendering Techniques '96*. Porto, Portugal: Springer-Verlag, 1996, pp. 21–30. ISBN: 3-211-82883-4. URL: <http://dl.acm.org/citation.cfm?id=275458.275461> (cit. on p. 3).
- [Kaj86] James T. Kajiya. “The Rendering Equation.” In: *SIGGRAPH Comput. Graph.* 20.4 (Aug. 1986), pp. 143–150. ISSN: 0097-8930. DOI: [10.1145/15886.15902](https://doi.org/10.1145/15886.15902). URL: <http://doi.acm.org/10.1145/15886.15902> (cit. on p. 3).

- [Kar12] Tero Karras. “Maximizing Parallelism in the Construction of BVHs, Octrees, and K-d Trees.” In: *Proceedings of the Fourth ACM SIGGRAPH / Eurographics Conference on High-Performance Graphics*. EGGH-HPG’12. Paris, France: Eurographics Association, 2012, pp. 33–37. ISBN: 978-3-905674-41-5. DOI: [10.2312/EGGH/HPG12/033-037](https://doi.org/10.2312/EGGH/HPG12/033-037). URL: <https://doi.org/10.2312/EGGH/HPG12/033-037> (cit. on p. 11).
- [Kel+19] Alexander Keller et al. “Are We Done with Ray Tracing?” In: *ACM SIGGRAPH 2019 Courses*. SIGGRAPH ’19. Los Angeles, California: ACM, 2019, 3:1–3:381. ISBN: 978-1-4503-6307-5. DOI: [10.1145/3305366.3329896](https://doi.org/10.1145/3305366.3329896). URL: <http://doi.acm.org/10.1145/3305366.3329896> (cit. on p. 3).
- [Kel97] Alexander Keller. “Instant Radiosity.” In: *Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques*. SIGGRAPH ’97. New York, NY, USA: ACM Press/Addison-Wesley Publishing Co., 1997, pp. 49–56. ISBN: 0-89791-896-7. DOI: [10.1145/258734.258769](https://doi.org/10.1145/258734.258769). URL: <https://doi.org/10.1145/258734.258769> (cit. on pp. 3, 19).
- [LWC12] Tzu-Mao Li, Yu-Ting Wu, and Yung-Yu Chuang. “SURE-based Optimization for Adaptive Sampling and Reconstruction.” In: *ACM Trans. Graph.* 31.6 (Nov. 2012), 194:1–194:9. ISSN: 0730-0301. DOI: [10.1145/2366145.2366213](https://doi.org/10.1145/2366145.2366213). URL: <http://doi.acm.org/10.1145/2366145.2366213> (cit. on p. 5).
- [LY19] Daqi Lin and Cem Yuksel. “Real-Time Rendering with Lighting Grid Hierarchy.” In: *Proc. ACM Comput. Graph. Interact. Tech. (Proceedings of I3D 2019)* 2.1 (2019), 8:1–8:17. DOI: [10.1145/3321361](https://doi.org/10.1145/3321361). URL: <http://doi.acm.org/10.1145/3321361> (cit. on pp. 3, 21).
- [NVI18] NVIDIA Corporation. *NVIDIA Turing GPU Architecture*. 2018. URL: <https://www.nvidia.com/content/dam/en-zz/Solutions/design-visualization/technologies/turing-architecture/NVIDIA-Turing-Architecture-Whitepaper.pdf> (cit. on p. 3).
- [OP11] Jiawei Ou and Fabio Pellacini. “LightSlice: Matrix Slice Sampling for the Many-lights Problem.” In: *ACM Trans. Graph.* 30.6 (Dec. 2011), 179:1–179:8. ISSN: 0730-0301. DOI: [10.1145/2070781](https://doi.org/10.1145/2070781).

## Bibliography

2024213. URL: <http://doi.acm.org/10.1145/2070781.2024213> (cit. on p. 4).
- [Par+10] Steven G. Parker et al. "OptiX: A General Purpose Ray Tracing Engine." In: *ACM Trans. Graph.* 29.4 (July 2010), 66:1–66:13. ISSN: 0730-0301. DOI: 10.1145/1778765.1778803. URL: <http://doi.acm.org/10.1145/1778765.1778803> (cit. on p. 3).
- [Rit+08] T. Ritschel et al. "Imperfect Shadow Maps for Efficient Computation of Indirect Illumination." In: *ACM Trans. Graph.* 27.5 (Dec. 2008), 129:1–129:8. ISSN: 0730-0301. DOI: 10.1145/1409060.1409082. URL: <http://doi.acm.org/10.1145/1409060.1409082> (cit. on p. 3).
- [Sch+17] Christoph Schied et al. "Spatiotemporal Variance-guided Filtering: Real-time Reconstruction for Path-traced Global Illumination." In: *Proceedings of High Performance Graphics*. HPG '17. Los Angeles, California: ACM, 2017, 2:1–2:12. ISBN: 978-1-4503-5101-0. DOI: 10.1145/3105762.3105770. URL: <http://doi.acm.org/10.1145/3105762.3105770> (cit. on pp. 5, 20).
- [SIP07] B. Segovia, J.C. Iehl, and B. Peroche. "Metropolis Instant Radiosity." In: *Computer Graphics Forum* (2007). ISSN: 1467-8659. DOI: 10.1111/j.1467-8659.2007.01065.x (cit. on p. 3).
- [SPD18] Christoph Schied, Christoph Peters, and Carsten Dachsbacher. "Gradient Estimation for Real-time Adaptive Temporal Filtering." In: *Proc. ACM Comput. Graph. Interact. Tech.* 1.2 (Aug. 2018), 24:1–24:16. ISSN: 2577-6193. DOI: 10.1145/3233301. URL: <http://doi.acm.org/10.1145/3233301> (cit. on pp. 5, 20).
- [TR98] Gabriel Taubin and Jarek Rossignac. "Geometric Compression Through Topological Surgery." In: *ACM Trans. Graph.* 17.2 (Apr. 1998), pp. 84–115. ISSN: 0730-0301. DOI: 10.1145/274363.274365. URL: <http://doi.acm.org/10.1145/274363.274365> (cit. on p. 11).
- [VG97] Eric Veach and Leonidas J. Guibas. "Metropolis Light Transport." In: *Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques*. SIGGRAPH '97. New York,

- NY, USA: ACM Press/Addison-Wesley Publishing Co., 1997, pp. 65–76. ISBN: 0-89791-896-7. DOI: [10.1145/258734.258775](https://doi.org/10.1145/258734.258775). URL: <https://doi.org/10.1145/258734.258775> (cit. on p. 3).
- [Wal+05] Bruce Walter et al. “Lightcuts: A Scalable Approach to Illumination.” In: *ACM Trans. Graph.* 24.3 (July 2005), pp. 1098–1107. ISSN: 0730-0301. DOI: [10.1145/1073204.1073318](https://doi.org/10.1145/1073204.1073318). URL: <http://doi.acm.org/10.1145/1073204.1073318> (cit. on pp. 4, 7, 16).
- [Wal+06] Bruce Walter et al. “Multidimensional Lightcuts.” In: *ACM Trans. Graph.* 25.3 (July 2006), pp. 1081–1088. ISSN: 0730-0301. DOI: [10.1145/1141911.1141997](https://doi.org/10.1145/1141911.1141997). URL: <http://doi.acm.org/10.1145/1141911.1141997> (cit. on p. 4).
- [Yuk19] Cem Yuksel. “Stochastic Lightcuts.” In: *High-Performance Graphics (HPG 2019)*. Strasbourg, France: The Eurographics Association, 2019. ISBN: 978-3-03868-092-5. DOI: [10.2312/hpg.20191192](https://doi.org/10.2312/hpg.20191192) (cit. on pp. 2, 4, 7).