

Slope-Space Integrals for Specular Next Event Estimation

GUILLAUME LOUBET, École Polytechnique Fédérale de Lausanne (EPFL)

TIZIAN ZELTNER, École Polytechnique Fédérale de Lausanne (EPFL)

NICOLAS HOLZSCHUCH, Inria, Univ. Grenoble-Alpes, CNRS, LJK

WENZEL JAKOB, École Polytechnique Fédérale de Lausanne (EPFL)



(a) Caustics created by a small light source and near-specular reflection.



(b) Underwater caustics due to sunlight and near-specular refraction.

Fig. 1. We address the problem of importance sampling light paths involving specular or near-specular reflection (a) or refraction (b). Our approach finds triangles that reflect or refract light from a point on a light source towards a particular shading point in the scene. Our technique can be used in addition to standard next event estimation and handles specular-diffuse-specular (“SDS”) sub-paths that are a well-known failure case of standard uni- and bidirectional path tracers. The images were rendered in 30m (a) and 5m (b) using an unidirectional path tracer together with the proposed sampling technique.

Monte Carlo light transport simulations often lack robustness in scenes containing specular or near-specular materials. Widely used uni- and bidirectional sampling strategies tend to find light paths involving such materials with insufficient probability, producing unusable images that are contaminated by significant variance.

This article addresses the problem of sampling a light path connecting two given scene points via a single specular reflection or refraction, extending the range of scenes that can be robustly handled by unbiased path sampling techniques. Our technique enables efficient rendering of challenging transport phenomena caused by such paths, such as underwater caustics or caustics involving glossy metallic objects.

We derive analytic expressions that predict the total radiance due to a single reflective or refractive triangle with a microfacet BSDF and we show that this reduces to the well known Lambert boundary integral for irradiance. We subsequently show how this can be leveraged to efficiently sample connections on meshes comprised of vast numbers of triangles.

Authors’ addresses: Guillaume Loubet, École Polytechnique Fédérale de Lausanne (EPFL), g.loubet.research@gmail.com; Tizian Zeltner, École Polytechnique Fédérale de Lausanne (EPFL), tizian.zeltner@epfl.ch; Nicolas Holzschuch, Inria, Univ. Grenoble-Alpes, CNRS, LJK, nicolas.holzschuch@inria.fr; Wenzel Jakob, École Polytechnique Fédérale de Lausanne (EPFL), wenzel.jakob@epfl.ch.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM. 0730-0301/2020/12-ART239 \$15.00

<https://doi.org/10.1145/3414685.3417811>

Our derivation builds on the theory of off-center microfacets and involves integrals in the space of surface slopes.

Our approach straightforwardly applies to the related problem of rendering glints with high-resolution normal maps describing specular microstructure. Our formulation alleviates problems raised by singularities in filtering integrals and enables a generalization of previous work to perfectly specular materials. We also extend previous work to the case of GGX distributions and introduce new techniques to improve accuracy and performance.

CCS Concepts: • **Computing methodologies** → **Rendering**; **Ray tracing**.

Additional Key Words and Phrases: specular transport, SDS paths, microfacet theory, glints

ACM Reference Format:

Guillaume Loubet, Tizian Zeltner, Nicolas Holzschuch, and Wenzel Jakob. 2020. Slope-Space Integrals for Specular Next Event Estimation. *ACM Trans. Graph.* 39, 6, Article 239 (December 2020), 13 pages. <https://doi.org/10.1145/3414685.3417811>

1 INTRODUCTION

Physically-based rendering of scenes containing specular materials like metal, glass, water, or plastic requires robust sampling strategies that can find light paths involving specular or near-specular interactions. Specular paths that connect the camera and light sources subtending a small solid angle often lead to visually striking effects in rendered images. Caustics induced by the glossy objects in Fig. 1a, and the underwater scene in Fig. 1b are examples of such specular paths. Uni- and bidirectional path tracing algorithms [Lafortune and Willems 1993; Veach and Guibas 1994] are widely used yet




Fig. 2. We address the problem of importance sampling light paths connecting two points p_v and p_l with one intermediate (near-)specular reflection or refraction p_s . We derive the relative contribution of each triangle along with a sampling procedure to choose positions within triangles. This enables us to find paths carrying large amounts of energy, which happens where half vectors align with the shading normal.

notoriously bad at sampling some types of specular light paths such as specular-diffuse-specular (“SDS”) or *near-SDS* sub-paths. Paths carrying significant amounts of energy are found with an insufficient probability, causing these methods to have an unacceptably low rate of convergence.

We address the problem of importance sampling specular sub-paths with exactly one specular interaction as shown in Fig. 2: given two positions p_l and p_v , we propose a technique for sampling an intermediate specular reflection or refraction p_s that carries energy from one position to the other. This entails two main steps: we first choose a triangle on a specular object based on its relative contribution. Following this, we sample a position within the triangle—this is important because the contributing region can be much smaller than the triangle itself. We find solutions to both steps via integration in the space of microfacet slopes, building on the theory of off-center microfacet models [Dupuy 2015; Dupuy et al. 2013; Olano and Baker 2010]. We show that integrals of GGX distributions of slopes [Trowbridge and Reitz 1975; Walter et al. 2007] over triangular domains reduce to well-known closed-form expressions of irradiance from polygonal lights [Heitz 2017; Lambert 1760] and introduce an approximation concerning the variation of viewing and lighting directions within triangles that enables robust importance sampling and unbiased results for rough materials. In the case of ideally specular materials, we propose a simple solution that does not require iterative root-finding techniques [Hanika et al. 2015a; Walter et al. 2007]. Our approach builds on a hierarchical data structure to efficiently identify the triangles involved in indirect specular transport at render time.

Our approach is closely related to previous work on sampling specular paths in refractive media [Walter et al. 2009], which we improve and generalize to the case of rough materials and reflection. In contrast to path guiding approaches [Müller et al. 2017; Reibold et al. 2018; Vorba et al. 2014] and algorithms based on walks on

specular manifolds [Hanika et al. 2015a; Jakob 2013; Jakob and Marschner 2012], our approach achieves robust global exploration and ensures high-quality sampling of this type of specular sub-paths. Its main limitation is that it can only sample sub-paths with exactly one intermediate specular interaction, hence more complex specular chains must be filtered or sampled using other techniques.

Our method can also be applied to the related problem of rendering specular glints on geometry with high-frequency normal maps [Yan et al. 2014]. Our approach extends this work to the case of GGX distributions of microfacets and it reduces artifacts by accounting for surface curvature and varying half vectors inside texels. We also propose an improved pruning strategy for the case of rough materials. Our formulation alleviates problems raised by singularities in filtering integrals and enables a generalization to perfectly specular materials that runs about twice as fast as in the case of rough materials with tiny amounts of roughness, due to simpler expressions and a more drastic pruning of the search space.

An open source implementation of the methods presented in this paper is available at rgl.epfl.ch/publications/Loubet2020SlopeSpace.

2 RELATED WORK

Our approach builds on microfacet theory and is related to specialized sampling techniques for specular transport. We briefly review both areas in this section.

Specular path sampling. Scenes with small light sources and narrowly peaked BRDFs are traditional failure cases of various rendering algorithms, which has spurred the development of a number of specialized approaches. Mitchell and Hanrahan [1992] proposed an algorithm for rendering caustics due to an implicitly defined surface, relying on interval arithmetic and a Newton-like iteration to find solutions. Walter et al. [2009] proposed a related technique for arbitrary triangle meshes with interpolated shading normals. Their work focuses on smooth dielectrics and introduces a spatio-directional acceleration data structure to prune large parts of the search space. We generalize their approach to the case of reflection and refraction from rough conductors and dielectrics, and we propose a new pruning strategy that substantially accelerates sampling of specular paths.

Jakob and Marschner’s [2013; 2012] Manifold Exploration method samples specular light paths within the framework of Metropolis Light Transport [Veach and Guibas 1997]. They leverage information provided by the differential geometry of the manifold of specular paths to sample valid configurations in the neighborhood of a given path, leading to a “manifold walk” in path space. Their approach was later considerably improved for the case of rough materials [Hanika et al. 2015b; Kaplanyan et al. 2014].

Manifold Next Event Estimation (MNEE) proposed by Hanika et al. [2015a] imports these tools into the framework of classical (i.e. non-MCMC) Monte Carlo methods. Their algorithm—like ours—samples specular light paths connecting two points. Starting from a deterministically chosen invalid initialization, their method performs a manifold walk to bring the starting path into a valid physical configuration. MNEE handles specular and near-specular refraction events, but it can at most find a single specular path connecting two

given endpoints. This is sufficient to render simple smooth geometry (e.g. spheres), but becomes problematic when the geometry is more complex and many solutions exist. The recent work of Zeltner et al. [2020] introduces a specular path sampling routine based on manifold walks that addresses some of these limitations.

Guided path tracing. Path guiding approaches exploit information about previously generated light paths to improve later sampling steps. Various types of caches have been proposed, including 5-dimensional incident radiance fields [Herholz et al. 2016; Müller et al. 2017; Vorba et al. 2014] and sets of high variance paths [Reibold et al. 2018]. In the case of caustics, the radiance function is characterized by high-frequency variation in both its spatial and directional arguments, which means that guiding techniques require a high-resolution representation built from a large number of relevant light paths. In common cases involving small/distant light sources or complex specular geometry, such paths are sampled with an insufficient probability, and guiding becomes ineffective. High-resolution normal maps further exacerbate this type of problem. We compare our method to state-of-the-art guiding techniques to demonstrate these limitations.

Glint rendering. Specular microstructure, e.g. specified via high-resolution normal maps, tends to produce severe aliasing and variance in rendered images. This happens due to small sub-pixel highlights that standard techniques fail to sample efficiently. Filtering approaches [Bruneton and Neyret 2011; Dupuy et al. 2013; Olano and Baker 2010; Xu et al. 2017] resolve this lack of convergence but generally result in a loss of visual fidelity and realism. Yan et al. [2014] proposed an algorithm for evaluating the average distribution of normals within a given surface region, which enables efficient rendering of glinty appearance. Like Walter et al. [2009], they rely on spatio-directional hierarchies to prune large groups of texels when evaluating normal queries. We reformulate their approach in the space of microfacet slopes (Sec. 5), while additionally accounting for surface curvature and variations of view and light directions. We also introduce a simpler approach for of smooth reflectance models, as well as better pruning strategies for the rough case.

Yan et al. [2016] improved the performance of the previous approach by approximating normal distributions using 4D Gaussian mixtures with millions of elements. This trades higher storage footprint (several gigabytes of memory for a typical normal map) in exchange for a roughly hundred-fold performance improvement. Our spatial-angular slope space hierarchies could also benefit from this optimization. Many extensions and variants of glint rendering techniques have been proposed in prior work, e.g., to support multiple scattering [Chermain et al. 2019; Raymond et al. 2016], directional product integrals [Gamboa et al. 2018], procedural glints [Jakob et al. 2014; Kuznetsov et al. 2019], wave optics [Werner et al. 2017; Yan et al. 2018] and real-time rendering [Zirr and Kaplanyan 2016]. Some of these methods have been formulated in slope space [Chermain et al. 2019; Zirr and Kaplanyan 2016].

Off-center microfacet models. We build on microfacet models with off-center normal distributions, meaning that their average normal differs from the geometric normal. Such distributions were used by Olano and Baker [2010] to filter normal maps, and Dupuy, Heitz

D	Distribution of normals
P	Normalized distribution of slopes
$\mathbf{p}_l, \mathbf{p}_v, \mathbf{p}_s$	3D points
$\mathbf{l}, \mathbf{v}, \mathbf{n}, \mathbf{h}$	3D vectors and normals
$\tilde{\mathbf{n}}, \tilde{\mathbf{h}}$	2D positions in slope space
Δ_i	Triangle in world space
$\tilde{\Delta}_i$	Triangle in slope space
f_s, D_s, P_s	BSDF and distributions at the point $\mathbf{p}_s$

Table 1. List of important notations.

et al. [2013] introduced the missing masking and shadowing functions. Schussler et al. [2017] recently added vertical microfacets to off-center distributions to ensure energy conservation on normal-mapped surfaces. A symmetric off-center BRDF model can be found in the work of Dupuy [2015]. In the supplemental material of the article Dong et al. [2015], Walter et al. [2015] introduced an off-center microfacet distribution of normals called the Ellipsoid NDF along with evaluation, sampling and shadowing-masking terms. This model is exactly equivalent to off-center GGX distributions as discussed in our supplemental material.

Integrals over normals and slopes. The spaces of microfacet slopes and projected directions appear in numerous prior works [Dupuy et al. 2013; Hanika et al. 2015b; Heitz 2014; Kaplanyan et al. 2014; Olano and Baker 2010; Smith 1967] and have proven to be convenient spaces for addressing light transport and filtering problems. In particular, Kaplanyan et al. [2016] rely on slope-space integrals for filtering specular highlights on specular surfaces with high curvature. Their approach had a limited accuracy in the case of GGX distributions because of the lack of analytic integrals of this distribution in an arbitrary polygonal domain. We provide these missing formulas in section 3.

3 SLOPE-SPACE INTEGRALS FOR SPECULAR TRIANGLES

In this section, we derive expressions that approximate the total amount of transport between two points $\mathbf{p}_v$ and $\mathbf{p}_l$ by means of a triangle with shading normals defined at its vertices and a specular or near-specular BSDF, as shown in Fig. 2. Such expressions enable the importance sampling of specular triangles when rendering caustics (Section 4), or specular texels when rendering glints (Section 5). We also discuss importance sampling techniques for the positions of specular interactions within a triangle.

As spatially varying roughness within triangles would lead to intractable integration problems in the general case, we assume that BSDF roughness is constant in each triangle. For the sake of simplicity, also we assume that the endpoints $\mathbf{p}_v$ and $\mathbf{p}_l$ are located on scene surfaces. Our results easily generalize to other cases, e.g., when $\mathbf{p}_l$ is a directional environment light source, or when $\mathbf{p}_v$ is a camera.

3.1 Integral of the radiance at $\mathbf{p}_v$

Given a position $\mathbf{p}_l$ (e.g. a position on a light source), the radiance due to paths with two segments connecting $\mathbf{p}_v$ and $\mathbf{p}_l$ by means of

a specular triangle Δ_i is given by an integral:

$$\int_{\Delta_i} f_v(\mathbf{o}_v, -\mathbf{v}) \frac{|\mathbf{v} \cdot \mathbf{n}_v| V(\mathbf{p}_v \leftrightarrow \mathbf{p}_s) |\mathbf{v} \cdot \mathbf{n}_s|}{\|\mathbf{p}_v - \mathbf{p}_s\|^2} f_s(\mathbf{v}, \mathbf{l}) |\mathbf{l} \cdot \mathbf{n}_s| L(\mathbf{p}_s, -\mathbf{l}) d\mathbf{p}_s \quad (1)$$

where $\mathbf{p}_s \in \Delta_i$, $\mathbf{l} = \overrightarrow{\mathbf{p}_s \mathbf{p}_1}$, $\mathbf{v} = \overrightarrow{\mathbf{p}_s \mathbf{p}_v}$, f_v and $\mathbf{n}_v$ are the BSDF and normal at $\mathbf{p}_v$, f_s and $\mathbf{n}_s$ are the BSDF and normal at $\mathbf{p}_s$, and $V(\mathbf{p}_v \leftrightarrow \mathbf{p}_s)$ is a visibility term whose value is either 0 or 1. The term $L(\mathbf{p}_s, -\mathbf{l})$ denotes the incident radiance at $\mathbf{p}_s$ from $\mathbf{p}_1$. For simplicity, we omit the direction $\mathbf{o}_v$ in the following equations.

Eq. 1 generally lacks an analytic solution that would be needed to make caustics and glints rendering practical. The difficulty in obtaining closed-form results is related to three challenges:

- View and light directions $\mathbf{v}$ and $\mathbf{l}$ vary non-linearly and affect most of the terms including the BSDF $f_s(\mathbf{v}, \mathbf{l})$ and the incident radiance $L(\mathbf{p}_s, -\mathbf{l})$.
- Variations in the shading normal effectively induce spatial variation in the BSDF of the specular surface.
- When f_s is a smooth conductor or a smooth dielectric, this integral reduces to a discrete sum of light path contributions. This is the approach studied by Walter et al. [2009] for dielectric materials. Even here, there is no explicit solution for valid light paths, and previous works [Hanika et al. 2015a; Walter et al. 2009] require iterative nonlinear solvers.

We will introduce simplifying assumptions to resolve these mathematical difficulties.

3.2 Far-field approximation

Far-field approximations neglect small variations in directions and/or positions involving a distant surface. They have been applied in the context of glint rendering [Jakob et al. 2014; Yan et al. 2014, 2016], where the simplified problem reduces to evaluating the averaged distribution of normals in a surface region.


We instead introduce a far-field approximation within each triangle Δ_i , which is very accurate for most combinations of shading points, light samples and triangles in a typical scene, in particular for high-quality meshes with small triangles. In practice, we ensure that this approximation always holds by using on-the-fly adaptive triangle subdivisions as discussed in section 4.6.

Assuming that $\mathbf{p}_v$ and $\mathbf{p}_1$ are distant with respect to triangle Δ_i , and that the variations of the BRDF at $\mathbf{p}_v$ and the incident light at $\mathbf{p}_s$ are negligible, most of the terms of Eq. 1 can be moved outside the integral. The problem then reduces to estimating the average BRDF over Δ_i for fixed directions $\mathbf{l}$ and $\mathbf{v}$:

$$\int_{\Delta_i} f_v(-\mathbf{v}) \frac{|\mathbf{v} \cdot \mathbf{n}_v| V(\mathbf{p}_v \leftrightarrow \mathbf{p}_s) |\mathbf{v} \cdot \mathbf{n}_s|}{\|\mathbf{p}_v - \mathbf{p}_s\|^2} f_s(\mathbf{v}, \mathbf{l}) |\mathbf{l} \cdot \mathbf{n}_s| L(\mathbf{p}_s, -\mathbf{l}) d\mathbf{p}_s \quad (2)$$

$$\approx f_v(-\mathbf{v}) \frac{|\mathbf{v} \cdot \mathbf{n}_v| V(\mathbf{p}_v \leftrightarrow \mathbf{p}_s) |\mathbf{v} \cdot \mathbf{n}_s|}{\|\mathbf{p}_v - \mathbf{p}_s\|^2} |\mathbf{l} \cdot \mathbf{n}_s| L(\mathbf{p}_s, -\mathbf{l}) \int_{\Delta_i} f_s(\mathbf{v}, \mathbf{l}) d\mathbf{p}_s. \quad (3)$$

We note that f_s remains spatially varying due to interpolated shading normals over Δ_i . When shading frames are used, Eq. 3 involves some form of spherical convolution with spherical triangles, for which there is no known closed-form expressions in general.



(a) Interpolation between 3 normals.

Fig. 3. Normals can be represented in the space of slopes, which can be visualized as projections of normals on the plane $z = 1$, shown in red. Barycentric interpolation between 3 normals can be computed either directly on Cartesian coordinates or in slope space. The support of the resulting distributions is the same spherical triangle (blue).

3.3 Off-center distributions of slopes

This problem simplifies when f_s is defined using a microfacet BSDFs with off-center microfacet distributions instead of shading frames. As in previous work, we define off-center distributions as *translations* of centered distributions in the space of microfacet slopes [Dupuy et al. 2013]. The conversion between normals and slopes defines a bijection:

$$\tilde{\mathbf{m}} = (\tilde{m}_x, \tilde{m}_y)^T = \left(-\frac{m_x}{m_z}, -\frac{m_y}{m_z} \right)^T \Leftrightarrow \mathbf{m} = \frac{(-\tilde{m}_x, -\tilde{m}_y, 1)^T}{\sqrt{\tilde{m}_x^2 + \tilde{m}_y^2 + 1}} \quad (4)$$

where $\tilde{m}_x$ and $\tilde{m}_y$ are the slopes of the microfacet with normal $\mathbf{m} = (m_x, m_y, m_z)^T$. As in previous work [Dupuy et al. 2013] we define and use normalized distributions in slope space:

$$D(\mathbf{m}) = \frac{P(\tilde{\mathbf{m}})}{|\mathbf{m} \cdot \mathbf{n}|^4}, \quad (5)$$

where P is the normalized distribution of slopes and $\mathbf{n}$ is the normal of the surface.

In the case of reflection, for instance, the integral term in Eq. 3 with such BSDF models yields:

$$\int_{\Delta_i} f_s(\mathbf{v}, \mathbf{l}) d\mathbf{p} = \int_{\Delta_i} \frac{D_s(\mathbf{h}) F(\mathbf{v}, \mathbf{h}) G(\mathbf{v}, \mathbf{l}, \mathbf{h})}{4|\mathbf{v} \cdot \mathbf{n}_s| |\mathbf{l} \cdot \mathbf{n}_s|} d\mathbf{p}_s \quad (6)$$

$$= \int_{\Delta_i} \frac{P_s(\tilde{\mathbf{h}}) F(\mathbf{v}, \mathbf{h}) G(\mathbf{v}, \mathbf{l}, \mathbf{h})}{4|\mathbf{v} \cdot \mathbf{n}_s| |\mathbf{l} \cdot \mathbf{n}_s| |\mathbf{h} \cdot \mathbf{n}_s|^4} d\mathbf{p}_s \quad (7)$$

$$= \frac{F(\mathbf{v}, \mathbf{h}) G(\mathbf{v}, \mathbf{l}, \mathbf{h})}{4|\mathbf{v} \cdot \mathbf{n}_s| |\mathbf{l} \cdot \mathbf{n}_s| |\mathbf{h} \cdot \mathbf{n}_s|^4} \int_{\Delta_i} P_s(\tilde{\mathbf{h}}) d\mathbf{p}_s \quad (8)$$

where P_s is a spatially-varying slope distribution, $\mathbf{h}$ denotes the half vector at $\mathbf{p}_s$ aligned with $\mathbf{v} + \mathbf{l}$, F is the Fresnel term and G is the masking/shadowing term [Cook and Torrance 1982; Walter et al. 2007].

3.4 Linear variations of slopes

The integrand P_s in Equation 8 is a spatially-varying distribution, which depends on the variations of the shading normals over the triangle Δ_i , that is, the variations of the mean slope of the distribution. This integral does not have a closed-form expression in general. However, we found that closed-form expressions can be derived




Fig. 4. Instead of neglecting variations of half-vectors in specular triangles, we use a linear approximation of such variations in slope space. This approximation is very accurate when points p_v and p_l are distant enough from the triangle (b), and it becomes less accurate when one of the point is close to the triangle, as half vector variations become increasingly non-linear (c). Note that points p_v and p_l are typically much farther from the triangle than what is shown case (b) (i.e., the approximation is superior), especially when working with high-resolution meshes and normal maps.

if the mean slope varies linearly over Δ_i . Therefore, we linearly interpolate the shading normals in the space of slopes:

$$\tilde{\mathbf{n}}(b_0, b_1, b_2) = b_0 \tilde{\mathbf{n}}_0 + b_1 \tilde{\mathbf{n}}_1 + b_2 \tilde{\mathbf{n}}_2 \quad (9)$$

where b_0 , b_1 and b_2 are the barycentric coordinates of p_s in Δ_i , $\tilde{\mathbf{n}}_0$, $\tilde{\mathbf{n}}_1$ and $\tilde{\mathbf{n}}_2$ are the slopes of the shading normals at the vertices of Δ_i , and $\tilde{\mathbf{n}}$ is the mean slope at p_s . We found that this interpolation in slope space is visually indistinguishable from interpolation of the Cartesian components in our experiments. As illustrated in Fig. 3, interpolating three normals in the spherical domain and in slope space results in spherical distributions of interpolated normals whose support is the same spherical triangle. This is similar to the approach of Yan et al. [2014] who interpolate orthogonal projections of the normals onto the tangent plane for the same reasons.

Moreover, we can relax the far-field assumption in Equation 8 and take into account first-order variations of half vector inside the triangle. This is similar to the approximations used in the work of Kaplanyan et al. [2016]—they use first-order approximations in projected pixel footprints onto the surface, while we use first-order approximations inside triangles Δ_i . We first compute the exact half vectors at the corners of a triangle and their corresponding slopes, and we replace the constant half-vector $\tilde{h}$ by a function $\tilde{h}_{\text{linear}}(b_0, b_1, b_2)$ that defines approximate slopes for half vectors in the triangle using linear interpolation:

$$\tilde{h}_{\text{linear}}(b_0, b_1, b_2) = b_0 \tilde{h}_0 + b_1 \tilde{h}_1 + b_2 \tilde{h}_2. \quad (10)$$

Examples of such linear approximations of the variations of half vectors can be found in Fig. 4. These approximations are accurate unless p_v or p_l lie within a small distance of the triangle. In practice, we always ensure that these approximations hold by subdividing triangles on the fly, as discussed in section 4.6.

Finally, we can write Equation 8 in the space of slopes. There is a bijection between positions p_s on the triangle Δ_i and mean

slopes $\tilde{\mathbf{n}}$, as well as a bijection between positions p_s and the approximate slopes $\tilde{h}_{\text{linear}}$ of the half-vector. Since these transformations are linear, the integral over Δ_i can be rewritten as an integral over a triangle in slope space. We now use the fact that the slopes $\tilde{\mathbf{n}}$ define a translation of a centered slope distribution P , and turn Eq. 8 into a slope-space integral over the triangle $\tilde{\Delta}_i = (\tilde{h}_0 - \tilde{n}_0, \tilde{h}_1 - \tilde{n}_1, \tilde{h}_2 - \tilde{n}_2)$:

$$\int_{\Delta_i} P_s(\tilde{h}_{\text{linear}}) d\mathbf{p}_s = \int_{\tilde{\Delta}_i} P(\tilde{\mathbf{s}} - \tilde{h}_{\text{linear}}) \left\| \frac{\partial \mathbf{p}_s}{\partial \tilde{\mathbf{s}}} \right\| d\tilde{\mathbf{s}} \quad (11)$$

$$= \left\| \frac{\partial \mathbf{p}_s}{\partial \tilde{\mathbf{s}}} \right\| \int_{\tilde{\Delta}_i} P(\tilde{\mathbf{s}} - \tilde{h}_{\text{linear}}) d\tilde{\mathbf{s}}. \quad (12)$$

The Jacobian term can be computed from the linear transformations from world space to barycentric coordinates and from barycentric coordinates to slope space, it remains constant for each particular triangle $\tilde{\Delta}_i$. The above equation shows that the combined assumptions of distant endpoints, off-center microfacets, and barycentric normal interpolation in slope space transform the total radiance due to an individual triangle into a slope-space integral over a triangular domain, as shown in Fig. 5a.

Note that these derivations are similar in spirit to the frameworks of Yan et al. [2014] (in the space of projected normals) and Kaplanyan et al. [2016] (based on differential geometry in surface filtering kernels).

3.5 Integrating slope distributions on arbitrary polygons

The integration of Beckmann slope distributions (i.e., 2D Gaussians) was previously studied by Yan et al. [2014]. Due to the lack of a closed-form expression for this integral, Yan et al. rely on a spline approximation.

We show that such integrals can be easily computed in the case of GGX distributions. First, the integration of arbitrary GGX distributions in polygonal domains (Fig. 5a) can be reduced to the case of the isotropic GGX distribution with unit roughness P_{std} through a simple linear transformation of the space of slopes, as shown in Fig. 5b and discussed in our supplemental material. The corresponding GGX distribution of normals D_{std} is a uniform distribution on the hemisphere:

$$\int_{\Delta} P_{\text{std}}(\tilde{\mathbf{m}}) d\tilde{\mathbf{m}} = \int_{\Delta} D_{\text{std}}(\mathbf{m}) |\mathbf{m} \cdot \mathbf{n}| d\mathbf{m} = \frac{1}{\pi} \int_{\Delta} |\mathbf{m} \cdot \mathbf{n}| d\mathbf{m}. \quad (13)$$

We have seen that a triangular domain in the space of slopes maps to a spherical triangle (Fig. 3). This gives a geometric interpretation for Eq. 13: GGX integrals in polygonal domains are equivalent to integrating *cosine-weighted* spherical triangles, that is, vertically-projected areas of spherical triangles as shown in Fig. 5c. Such integrals have a simple analytic expression:

$$\frac{1}{\pi} \int_{\Delta} |\mathbf{m} \cdot \mathbf{n}| d\mathbf{m} = \frac{1}{2\pi} \sum_{i=0}^2 \text{acos}(\mathbf{n}_i \cdot \mathbf{n}_j) \left(\frac{\mathbf{n}_i \times \mathbf{n}_j}{\|\mathbf{n}_i \times \mathbf{n}_j\|} \cdot \mathbf{n} \right). \quad (14)$$

where $\mathbf{n}_i$ are directions towards the vertices of Δ . This is a classical result from the work of Lambert [1760] which had various applications in Computer Graphics, in particular for rendering polygonal area lights [Heitz et al. 2016; Nishita and Nakamae 1985]. A complete derivation can be found in the work of Heitz [2017] and more details




Fig. 5. We show that the integral of GGX distributions of slopes P in triangular domains (a) have simple closed-form expressions. Arbitrary cases reduce to the case of an isotropic distribution with unit roughness (b) through a simple linear transformation. Then, the integral corresponds to the vertically-projected area of a spherical triangle (c), which has a simple analytic expression. We also importance sample the distribution restricted to triangles (d). These mathematical tools are at the core of our algorithms for rendering caustics and high-resolution normal maps.

can be found in our supplemental material. We note that computing these analytic expressions is not consistently faster than the spline approximations for Gaussian integrals from Yan et al. [2014], depending on the input parameters.

3.6 Sampling slopes within triangles

While the previously used slope integral could be used as an efficient radiance estimator, doing so would introduce bias in the rendering algorithm. To use our technique for unbiased rendering, we must therefore sample a position within the triangle and evaluate the associated visibility terms and BSDFs exactly.

Uniform sampling in triangles would be an exceedingly poor sampling strategy because the energy is often concentrated in small subregions of the triangle, where half vectors align with interpolated normals. A better solution is to importance sample positions proportionally to the integrand. This can be done in slope space, and it reduces to importance sampling a slope in the distribution P restricted to the triangle $\tilde{\Delta}_i$ as shown in Fig. 5d. Since there is a bijection from $\tilde{\Delta}_i$ and the geometric triangle Δ_i , the sampled slope corresponds to a particular light path with known density.

We have seen that an equivalence between GGX integrals and cosine-weighted spherical triangles can be leveraged for deriving closed-form integrals in polygonal domains (Sec. 3.5). For the same reasons, importance sampling a GGX distribution restricted to a triangle in the space of slopes is equivalent to uniform sampling of vertically projected spherical triangles, as shown in Fig. 5c. While exact procedures exist for projected spherical caps [Peters and Dachsbacher 2019; Ureña and Georgiev 2018] there is no known analytic solution that achieves stratified sampling for projected spherical triangles. Ureña et al. [2000] proposed an approximate importance sampling procedure based on adaptive subdivisions of the sampling domain, while Arvo [2001] relies on the numerical inversion of a cumulative distribution function. We rely on a similar approach and provide additional details in our supplemental material. Note that this problem is distinct from the problem of sampling of spherical triangles, which has a known solution [Arvo 1995]. In practice, variation of the cosine term within triangles is often very small, in

which case this sampling technique provides a useful alternative approach.

3.7 Singular distributions of slopes

The derivations in this section are valid whenever integration in slope space is feasible, which requires a bijective mapping between positions in world and slope space. This means that the triangle in slope space must not be degenerate (e.g. collapsed to a line or point). In the far-field approximation, this would be a problem, since flat or cylinder-like geometry would result in exactly such degenerate cases. During this project, we initially derived specialized integrals and sampling procedures to handle degeneracies, which introduced considerable additional complexity.

However, since the half vector typically varies across the triangle, and since we correct the slope distribution accordingly, the slope-space triangles $\tilde{\Delta}_i$ are generally not singular. Some singularities could remain for scenes with both an orthographic camera and a directional emitter such as an environment map, or when the variation of slopes due to interpolated vertex normals exactly compensate the variation of slopes of the half vector, which we have not observed in our experiments.

3.8 The case of smooth conductors and dielectrics

So far, our discussion has been focused on rough conductors and dielectrics. Applying our approximations to the case of smooth materials with Dirac delta slope distributions reduces the slope-space integrals to either zero or one specular light path per triangle. The light path corresponds to the solution of the equation

$$\tilde{\mathbf{n}}(b_0, b_1, b_2) = \tilde{\mathbf{h}}_{\text{linear}}(b_0, b_1, b_2), \quad (15)$$

which can be solved by inverting a 2×2 matrix. The triangle contains a solution if the resulting barycentric coordinates define a point inside Δ_i . Eq. 11 reduces to

$$\int_{\Delta_i} P_s(\tilde{\mathbf{h}}) d\mathbf{p}_s = \int_{\tilde{\Delta}_i} \delta_{\tilde{\mathbf{h}}}(\tilde{\mathbf{s}}) \left\| \frac{d\mathbf{p}_s}{d\tilde{\mathbf{s}}} \right\| d\tilde{\mathbf{s}} = \left\| \frac{d\mathbf{p}_s}{d\tilde{\mathbf{s}}} \right\|, \quad (16)$$

which uses the fact the $\tilde{\Delta}_i$ is a non-singular distribution of slopes defining a bijection from surface to slopes. Although this approach is numerically stable, it is biased and it can lead to artifacts when our far-field approximation does not hold. As in previous work [Walter et al. 2009] and as discussed in section 4.6, we selectively subdivide triangles in such difficult cases. The subdivision criterion allows for controlling the error and its role is similar to error thresholds in numerical root finding techniques based on Newton iterations [Hanika et al. 2015a; Jakob and Marschner 2012; Walter et al. 2009].

4 APPLICATION TO CAUSTIC RENDERING

The derivations of section 3 provide an algorithm to connect two given scene points via an intermediate specular vertex. We now show how this algorithm can be used as an additional Next Event Estimation (NEE) strategy in path tracing-style algorithms—we refer to this approach as *Specular Next Event Estimation* (SNEE).

4.1 Overview

We add SNEE to the core loop of a standard unidirectional path tracer. Like classical NEE, which samples direct connections to light




Fig. 6. We precompute spatial bounds of reflected and refracted caustics due to specular triangles and build a pruning hierarchy that supports efficient search of specular sub-paths. We also store expected irradiance values to enable adaptive stochastic pruning of specular geometry.

sources, SNEE produces sub-paths that connect to light sources, and the corresponding path throughputs and probability densities can be computed to combine all available sampling strategies using multiple importance sampling.

SNEE requires a point $\mathbf{p}_l$ on an emitter that can be generated using any available strategy. We then sample an interaction $\mathbf{p}_s$ on the (near-)specular surface. The contribution of the resulting path is given by

$$f_v(-\mathbf{v})|\mathbf{v} \cdot \mathbf{n}_v| \frac{|\mathbf{v} \cdot \mathbf{n}_s|}{\|\mathbf{p}_v - \mathbf{p}_s\|^2} f_s(\mathbf{v}, \mathbf{l}) |\mathbf{l} \cdot \mathbf{n}_s| L(\mathbf{p}_s, -\mathbf{l}) \frac{1}{p_s p_\Delta p_l}, \quad (17)$$

where p_s is the area density of $\mathbf{p}_s$ within the triangle, p_Δ is the discrete probability of choosing triangle Δ_i (i.e., proportional to its predicted contribution), and p_l is the probability for $\mathbf{p}_l$ with respect to the *solid angle measure* at $\mathbf{p}_s$. The visibility terms $V(\mathbf{p}_v \leftrightarrow \mathbf{p}_s)$ and $V(\mathbf{p}_s \leftrightarrow \mathbf{p}_l)$ are computed, and the path is rejected if occlusion is detected.

4.2 Hierarchical pruning of triangles

A naïve implementation of SNEE would first estimate the contributions of all specular triangles towards transport between two given points $\mathbf{p}_v$ and $\mathbf{p}_l$ and then importance sample one of the involved triangles. In practice, the vast majority of triangles do not contribute significantly towards this estimate. Walter et al. [2009] addressed this problem by precomputing a bounding volume hierarchy to prune large groups of triangles during the search for connection sub-paths. Each node of the BVH stores an axis-aligned bounding box and cone that bound the positions and shading normals of geometry in the associated subtree. To achieve good performance, a suitable pruning criterion given two scene points $\mathbf{p}_v$ and $\mathbf{p}_l$ is key. Walter et al. [2009] rely on a spindle test to efficiently prune BVH nodes using the property that a single refraction can only bend the light by a limited amount, as well as other tests that entail the computation of a bounding cone of half vectors for each visited node during a traversal of the BVH.

This pruning strategy was specially designed for the refractive case, and its generalization to the case of specular reflection remains problematic: the set of valid specular reflection configurations is not restricted to a spindle, and there is no equivalent test for pruning triangles efficiently based on their position with respect to $\mathbf{p}_v$ and $\mathbf{p}_l$. Reflective objects tend to scatter light towards the entirety of the scene, which means that essentially all scene surfaces receive

low-energy caustics. Although they barely contribute to the pixel value, such caustics still produce considerable variance when standard BSDF and emitter sampling techniques are used. At the same time, this means that pruning becomes largely ineffective, as range queries will tend to find many candidate triangles with a valid specular configuration. Finally, the method also lacks awareness about global visibility, i.e., it finds triangles with the right orientation even if occlusion prevent them from being visible from $\mathbf{p}_v$, $\mathbf{p}_l$, or both. We previously experimented with combinations of SNEE with the pruning hierarchy of Walter et al. [2009] and found that the drawbacks listed above make such an approach unnecessarily costly for the case of specular reflection.

We address all these problems by precomputing a different type of hierarchical data structure that enables fast queries for candidate triangles, whose contribution we subsequently evaluate using SNEE.

The data structure is generated in a brief pre-rendering step, which entails tracing rays from the light sources towards uniformly sampled positions in each specular triangle. We reflect and refract these rays and compute an axis-aligned bounding box containing all intersections with the non-specular scene geometry (Fig. 6) together with the associated amount of energy. In the case of refractive geometry, we separately bound reflected and refracted rays. Next, the per-triangle bounds are assembled into a weighted *bounding volume hierarchy* (BVH), where each node records the maximum expected irradiance from a specular triangle. Tree construction proceeds top-down, using the builder provided in PBRT [Pharr et al. 2016].

Following this step, finding specular triangles that participate in caustic light transport at a surface position reduces to a hierarchical range query. Additionally, the irradiance values stored within BVH nodes enable an efficient stochastic sampling procedure that focuses computation on salient parts of the domain in the case where many triangles contribute.

4.3 Sampling rate for the precomputation

One limitation of this approach is that it does not strictly guarantee that the precomputed bounding boxes contain all the challenging caustics, meaning that some specular light paths are not sampled with SNEE. Caustics that do not lie in the precomputed BVH are sampled using standard techniques to achieve unbiased results. However, it is crucial that the precomputed per-triangle bounds contain all the light paths that produce outliers using standard techniques.

A sufficient number of rays thus have to be traced in the pre-computation phase to account for difficult visibility, non-singular light sources and material roughness. In our implementation, we use 500 samples per triangle for smooth materials and up to 2k samples for rough materials to ensure a dense sampling of microfacet distributions. In practice, we found that this approach is robust—in particular, no filtering or post-processing of any kind was used in any of our results. Insufficient sampling rates result in residual outliers in rendered images, or a loss of energy if these are filtered.

4.4 Selective activation of SNEE

Since SNEE is more computationally demanding than other sampling techniques, it is essential that it is only used when truly



Fig. 7. Equal time rendering (2 minutes) of a unidirectional path tracer with our SNEE compared to other unbiased algorithms. Top: a challenging scene with smooth and rough conductors. Bottom: a slightly modified version of the swimming pool scene. We reduced the resolution of the water geometry from 1.8M to 120k triangles, and we use a thin dielectric approximation for the refraction through the windows. SNEE robustly finds SDS paths that produce outliers with standard BSDF and emitter sampling techniques. Metropolis Light Transport and Manifold Exploration (MLT+ME) efficiently explore some regions of the path space but images suffer from typical large-scale artifacts and incorrect brightness. We also compare SNEE with the improved version of the Practical Path Guiding (PPG) algorithm [Müller 2019; Müller et al. 2017]. This algorithm performs well in the case of refraction but offers little benefits compared to standard path tracing for reflection, due to the higher frequency of the light transport in the spatial and directional domain for such caustics. SNEE produces low-variance estimates for sub-paths with one specular interaction but it does not sample longer specular chains, which explains the remaining outliers, *e.g.*, inter-reflection between the vases (top) or between the water and the windows (bottom). The rendering time includes training passes for PPG and the precomputation passes for SNEE, which took 40s (top) and 20s (bottom).

needed. The per-triangle bounds should be as compact as possible, and they should be restricted to those paths that add significant variance when using standard sampling techniques. Construction of the BVH therefore ignores paths, whose luminance using standard BSDF/emitter sampling techniques falls below a user-specified threshold. This parameter is intuitive: it is a bound of the maximum desired value of outlier samples. This strategy enables targeted application of SNEE to truly challenging specular geometry and light sources. When the illumination is represented using an environment map, this strategy can for instance restrict SNEE to directions that map to the sun or other bright light sources encoded in the environment map image. Standard sampling techniques are used for all the specular sub-paths that are not sampled with SNEE to ensure unbiasedness of the resulting estimator.

4.5 Stochastic pruning at render time

We use our pruning data structure to query candidate specular triangles for a shading point at render time. Then, we importance sample one triangle based on its predicted contribution as well as a specular interaction inside the triangle (Sec. 3). Our method achieves nearly-perfect importance sampling of specular light paths if all the candidate triangles are queried. This is also the approach of Walter et al. [2009]. However, this requires estimating the contribution of all the candidate triangles independently of their contribution to the image. This is generally suboptimal, in particular in scenes with low-energy caustics as discussed in Sec. 4.2. This can be improved by stochastically pruning nodes during the traversal of the BVH. We define the probability of exploring a node using

$$\text{pdf}_{\text{node}} = \min\left(1, \frac{I_{\text{max}}}{T}\right), \quad (18)$$

where $I_{\max}$ is the precomputed maximum expected irradiance of the node, and T is a user-specified threshold. This effectively enables adaptive stochastic pruning while preventing weighted path contributions from being excessively large.

4.6 Triangle subdivision

Our derivations in section 3 rely on far-field assumptions that are only accurate when $\mathbf{p}_v$ and $\mathbf{p}_l$ are sufficiently distant from the triangle (relative to the scale of the triangle itself). In cases where one of the points (typically $\mathbf{p}_v$) is close to the triangle, we recursively subdivide the triangle until the distance between the points and the sub-triangles is N times larger than the size of the sub-triangles. We use values of N between 5 and 20 depending on the roughness of the material (a higher precision is needed for small values of α). We also define a maximum depth of 20 for the subdivisions and ignore the sub-triangles when this maximum is reached. In practice, for high resolution meshes and distant lights, subdivision only occurs in a small number of pixels, e.g., when a specular object directly placed onto a floor that receives caustics. In general, the subdivision does not have a strong effect on overall runtime performance.

4.7 Implementation and results

We implemented our algorithm in Mitsuba 2 [Nimier-David et al. 2019]. All images in this article were rendered on an Intel i9-7960X workstation.

Equal time comparisons. Fig. 7 compares SNEE to other unbiased rendering algorithms on difficult scenes with specular geometry and small light sources. SNEE robustly samples SDS paths that produce outliers in both uni- and bidirectional path tracing. SNEE also outperforms the guided path tracer introduced by Müller et al. [2017] for reflected caustics. Robustly guiding light paths for such caustics requires very high resolution spatial and directional data structures that are not robustly trained due to the tiny probability of sampling important light paths. However, this path guiding approach tends to perform better for refractive caustics due to a lower cost per path compared to SNEE, at the cost of a larger memory footprint. Metropolis Light Transport and Manifold Exploration (MLT+ME) achieve good local exploration but the images show incorrect brightness and rendering artifacts at these low sample counts.

Comparison with MNEE. Fig. 8 shows comparisons between Manifold Next Event Estimation (MNEE) and our approach. Both techniques provide an additional connection strategy in path tracers, but they behave very differently. MNEE only supports a unique and deterministic initialization for the manifold walk for each pair of endpoints, otherwise the probability of finding a particular path lacks a closed-form expression. As a consequence, MNEE can only sample one highlight for each pair of endpoints and it offers limited benefits for rendering challenging caustics. MNEE is also restricted to refraction since it is unclear what would be a good initialization for reflection. In contrast, SNEE finds all the important light paths for refraction and reflection, but it is restricted to sub-paths with exactly one specular interaction.




Fig. 8. Comparison between path tracing with manifold next event estimation (PT+MNEE) and our specular next event estimation (PT+SNEE). Manifold next event estimation works well in scenes with simple geometry (top) but lacks robustness when multiple light paths connect a shading point to a light source across a refractive boundary (bottom). Note that unlike MNEE, our approach is not specific to refraction and it also noticeably reduces variance caused by reflective caustics on the ground plane. The precomputation for SNEE takes 3 seconds for these scenes.




Fig. 9. Equal sample comparison (20 spp) between our pruning strategy and the pruning strategy proposed by Walter et al. [2009]. The previous strategy performs well for refractive caustics (left) but poorly for reflection. Our precomputed pruning hierarchy supports faster specular triangles queries and adaptive stochastic pruning that make SNEE practical for both types of specular sub-paths. Times do not include our precomputation passes, which took 15s (left) and 20s (right). Specular interactions inside selected triangles have been sampled using our techniques in all images.

Comparison with Walter et al. [2009]. We provide comparisons with the pruning strategy of Walter et al. [2009] in Fig. 9. Our approach is significantly faster due to the combined effects of stochastic sampling and our improved approach to hierarchical pruning.




Fig. 10. Various scenes rendered with equal sample counts for PT+SNEE (b, our), and equal time comparison against standard path tracing (a). These results show how the cost of SNEE increases with the roughness of the material and the complexity of the caustics. Images in (a) and (b) only contain contributions from light paths that can be sampled by SNEE.

The performance difference is particularly large in the case of reflective caustics. Our approach achieves similar levels of variance at equal sample counts despite the stochastic pruning of specular triangles.

Influence of roughness and geometry on performance. Additional comparisons with path tracing can be found in Fig. 10 for various roughness values. SNEE outperforms standard sampling techniques in all cases. The cost of SNEE queries depends on the number of triangles that significantly contribute to the light transport between the two query points, and this number increases with the roughness of the material and the amount of bumps on the surface. This can be observed in the images of Fig. 10, rendered at equal SNEE sample counts.

4.8 Discussion

The cost of SNEE queries depends on the number of triangles that contribute non-negligible energy. Therefore, SNEE scales well with the number of triangles in the scene if the number of connection paths remains reasonable. For instance, SNEE can efficiently sample the important light paths in underwater caustics like in Fig. 1b even if the ocean is modeled using a large-scale high-resolution mesh, since only a few specular paths contribute to each shading point. However, importance sampling indirect reflection from complex bumpy surfaces like the ones in Fig. 1a is less efficient, since many bumps redirect light towards a single point on the table. The cost of SNEE scales linearly with the number of light sources that produce caustics, since one acceleration structure must be computed for each of these lights.

5 SLOPE-SPACE INTEGRALS FOR HIGH-RESOLUTION NORMAL MAPS

Specular surfaces with high-resolution normal maps often result in highly realistic glinty appearance, but rendering of such surfaces using standard pixel sampling is impractical when the light source subtends a small solid angle. Glints are extremely bright sub-resolution highlights, whose size is typically much smaller than normal map texels—which themselves project to small fractions of a pixel.

Yan et al. [2014] addressed this problem by evaluating averaged normal distributions based on specified filter kernels on the specular surface. They derived their filtering equations using normals following orthogonal projection on the tangent plane, and they define the distribution of microfacets of the BRDF as a 2D Gaussian distribution in this space. Unfortunately, texel contributions are given by integrals that do not have closed-form expressions. In this section, we reformulate such filtering integrals for normal maps in slope space, and we show that the problem of importance sampling highlights in a normal map is very similar to the problem of importance sampling highlights on triangle meshes (Sec. 3). Our reformulation of previous work in the space of microfacet slopes brings several benefits that we discuss in this section. We note that transitioning from projected normals to surface slopes does not require fundamental changes to the hierarchical data structure used to implement NDF queries [Yan et al. 2014].

5.1 Filtering integrals for normal maps

The filtering integral for normal maps is given by

$$\int_S W(\mathbf{p}_s) f_s(\mathbf{v}, \mathbf{l}) |\mathbf{l} \cdot \mathbf{n}| d\mathbf{p}_s, \quad (19)$$

where W is a normalized filter kernel on the surface $\mathcal{S}$, and $f_s(\mathbf{v}, \mathbf{l})$ denotes BRDF at the point $\mathbf{p}_s$. The vectors $\mathbf{v}$ and $\mathbf{l}$ are the directions towards the camera ($\mathbf{p}_v$) and a sampled position on a light source ($\mathbf{p}_l$). A suitable filter W can be determined using ray differentials [Igehy 1999] or covariance tracing [Belcour et al. 2017]. As in the work of Yan et al. [2014], we assume that the roughness of the base material is not spatially varying, and that the normals of the map are interpolated within texture-space triangles. We can split Eq. 19 into sub-integrals over texture-space triangles:

$$\sum_{\Delta_i \in \mathcal{T}} \int_{\Delta_i} W(\mathbf{t}) f_s(\mathbf{v}, \mathbf{l}) |\mathbf{l} \cdot \mathbf{n}| \left\| \frac{\partial \mathbf{p}_s}{\partial \mathbf{t}} \right\| d\mathbf{t}, \quad (20)$$

where $\mathbf{t}$ is a position in the texture space $\mathcal{T}$, Δ_i is a triangle in the normal map and $\|\partial \mathbf{p}_s / \partial \mathbf{t}\|$ is the Jacobian determinant of the transformation between areas in world space and texture space. Since we use barycentric interpolation of normals in triangles Δ_i , there are bijections between texture coordinates and slopes that allow turning the integral in Eq. 20 into a slope-space integral:

$$\int_{\tilde{\Delta}_i} W(\tilde{\mathbf{s}}) f_s(\mathbf{v}, \mathbf{l}) |\mathbf{l} \cdot \mathbf{n}| \left\| \frac{\partial \mathbf{p}_s}{\partial \tilde{\mathbf{s}}} \right\| d\tilde{\mathbf{s}} \quad (21)$$

$$= \int_{\tilde{\Delta}_i} W(\tilde{\mathbf{s}}) P(\tilde{\mathbf{s}} - \tilde{\mathbf{h}}) \frac{F(\mathbf{v}, \mathbf{h}) G(\mathbf{v}, \mathbf{l}, \mathbf{h})}{4|\mathbf{v} \cdot \mathbf{n}| |\mathbf{h} \cdot \mathbf{n}|^4} \left\| \frac{\partial \mathbf{p}_s}{\partial \tilde{\mathbf{s}}} \right\| d\tilde{\mathbf{s}}. \quad (22)$$

This integral resembles the main result of section 3: it reduces to the evaluation of the contribution due to a single light path in the case of smooth conductors (Sec. 3.8), or it can be importance sampled when the material is rough (Sec. 3.6).

Compared to the work of Yan et al. [2014], our approach benefits from well-known masking-shadowing terms for GGX distributions and closed-form expressions for integrals over triangular domains. However, our framework cannot account for the spatially varying weight $W(\tilde{\mathbf{s}})$ in each triangle in the normal map. We found that is not a problem in practice since the triangles are much smaller than the filtering kernel (we are filtering high-resolution normal maps) and variations of $W(\tilde{\mathbf{s}})$ inside triangles can be neglected when importance sampling highlights.

5.2 Accounting for varying half vectors and curvature

We can apply the techniques described in sections 3.2 and 3.4 to improve the accuracy of the filtering integral, while at the same time avoiding potential issues due to singular slope distributions (Sec. 3.7). We correct the distributions in each normal map triangle Δ_i using the derivatives of the shading normal at the center of the filtering kernel:

$$\tilde{\mathbf{n}}_s(\mathbf{t}) \approx \frac{\partial \tilde{\mathbf{n}}}{\partial u}(\mathbf{t}_u - c_u) + \frac{\partial \tilde{\mathbf{n}}}{\partial v}(\mathbf{t}_v - c_v), \quad (23)$$

where c_u and c_v are the texture-space coordinates of the center of the filter kernel. We also estimate the half vector variations in texture space:

$$\tilde{\mathbf{h}}(\mathbf{t}) \approx \tilde{\mathbf{h}}_c + \frac{\partial \tilde{\mathbf{h}}}{\partial u}(\mathbf{t}_u - c_u) + \frac{\partial \tilde{\mathbf{h}}}{\partial v}(\mathbf{t}_v - c_v), \quad (24)$$

where $\tilde{\mathbf{h}}_c$ correspond to the slopes of the half vector at the center of the filter kernel. The contribution of the normal map triangle Δ_i




Fig. 11. Clamping normal distributions when evaluating rough glints results in a small loss of energy, and conservative clamping tends to decrease the performance of glint algorithms (top). We use more aggressive clamping (bottom) and account for the missing energy using an efficient sampling strategy. This approach significantly improves the performance of glint evaluation in the case of rough conductors and dielectrics.

can be computed by integrating the distribution of slopes between the points

$$\tilde{\mathbf{h}}(\mathbf{t}_j) - \tilde{\mathbf{n}}_s(\mathbf{t}_j) - \tilde{\mathbf{n}}_j, \quad j \in \{0, 1, 2\}$$

where $\mathbf{t}_j$ are the texture coordinates of the vertex j in Δ_i and $\tilde{\mathbf{n}}_i$ are the corresponding slopes from the normal map.

The supplemental material showcases several animated results created using our slope-space representation, while accounting for surface curvature. These results demonstrate the good stability of the highlights even in challenging cases with disabled normal interpolation or highly-curved geometry, which produce slope distributions that are almost singular.

5.3 Pruning tests for texels

As in previous work, we rely on a quad-tree hierarchy to prune large regions of the normal map when integrating the response over a particular filter kernel. Each tree node corresponds to a rectangular region of the normal map and specifies a 2D box that bounds the slopes of underlying texels. We modified pruning tests to account for surface curvature and half vector variations inside the filter footprint. In the case of rough materials, we only consider clamped distributions of slopes during the queries. We account for the missing energy by sampling positions in the filter kernel




Fig. 12. We extend previous work on glint rendering to the case of smooth conductors and dielectrics. Glints from smooth materials (top) are rendered approximately twice as fast as glints with small amounts of roughness (bottom) with our implementation, due to simplified pruning tests and evaluation. Images to the right show equal-time rendering without filtering, with a standard path tracer. All images are rendered at resolution 1024×1024 , with a normal map of resolution $8k \times 8k$.

and by keeping the contributions that correspond to the tail of the distribution. As shown in Fig. 11, this improves the performance of the queries compared to conservative truncation of the distribution, while ensuring that the entire distribution is correctly sampled.

5.4 Generalization to smooth materials

Our new and simpler solution that is specific to perfectly smooth materials (section 3.8) leads to an algorithm that runs about twice as fast as the case of rough materials with small amounts of roughness, while producing similar appearance as shown in Fig. 12. This speedup is a consequence of more aggressive pruning that becomes possible as bounding boxes do not have to be expanded to account for the intrinsic roughness of the material. The evaluation of the filtered BRDF is also faster since evaluating a highlight in a texel reduces to evaluating one Jacobian term.

5.5 Multiple importance sampling

We only enable the evaluation of filtered BRDFs for small lights, and rely on standard sampling techniques for smooth lights (e.g., environment maps). We use multiple importance sampling to robustly

render flat regions of the normal map [Yan et al. 2016]. Computing the probability density function of sampled directions for a particular filter kernel resembles evaluation of the filtered BRDF.

6 CONCLUSION

We introduced a new technique to sample specular reflection or refraction events connecting two given scene points, typically a shading point and a light source. Our method, named Specular Next Event Estimation, can supplement standard sampling techniques in path tracing-like algorithms and greatly improves the convergence of supported path classes, extending the scope of unidirectional rendering to scenes containing caustics and specular-diffuse-specular light paths. Building on a closed-form solution of a slope-space integral, our method supports surfaces ranging from ideally specular to near-specular and even highly rough materials, although standard techniques may be preferable in the latter case. Our first-order approximation of half vector variations largely eliminates problems raised by singularities in previous formulations of glint filtering, leading to a generalization to ideally specular materials that runs faster than the case of materials with tiny amounts of roughness. The main limitation of our method is its restriction to specular chains with a single vertex, hence other strategies remain necessary to handle specular transport with multiple interactions. Addressing this significantly more challenging case is an important topic for future work.

REFERENCES

- James Arvo. 1995. Stratified Sampling of Spherical Triangles. In *Proceedings of the 22nd Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '95)*. 2. <https://doi.org/10.1145/218380.218500>
- J. Arvo. 2001. Stratified sampling of 2-manifolds. In *ACM SIGGRAPH 2001 Courses*.
- Laurent Belcour, Ling-Qi Yan, Ravi Ramamoorthi, and Derek Nowrouzezahrai. 2017. Antialiasing Complex Global Illumination Effects in Path-space. *ACM Transactions on Graphics* 36, 1 (Jan. 2017). <https://doi.org/10.1145/2990495>
- Eric Bruneton and Fabrice Neyret. 2011. A Survey of Non-linear Pre-filtering Methods for Efficient and Accurate Surface Shading. *IEEE Transactions on Visualization and Computer Graphics* (2011). <http://maverick.inria.fr/Publications/2011/BN11>
- Xavier Chermant, Frédéric Clauss, and Stéphane Mérouille. 2019. Glint Rendering based on a Multiple-Scattering Patch BRDF. *Computer Graphics Forum* 38, 4 (2019), 27–37. <https://doi.org/10.1111/cgf.13767>
- R. L. Cook and K. E. Torrance. 1982. A Reflectance Model for Computer Graphics. *ACM Trans. Graph.* 1, 1 (Jan. 1982), 7–24. <https://doi.org/10.1145/357290.357293>
- Zhao Dong, Bruce Walter, Steve Marschner, and Donald P. Greenberg. 2015. Predicting Appearance from Measured Microgeometry of Metal Surfaces. *ACM Trans. Graph.* 35, 1 (2015), 9:1–9:13.
- Jonathan Dupuy. 2015. *Photorealistic Surface Rendering with Microfacet Theory*. Theses. Université Claude Bernard - Lyon I ; Université de Montréal. <https://hal.archives-ouvertes.fr/tel-01291974>
- Jonathan Dupuy, Eric Heitz, Jean-Claude Iehl, Pierre Poulin, Fabrice Neyret, and Victor Ostromoukhov. 2013. Linear Efficient Antialiased Displacement and Reflectance Mapping. *ACM Transactions on Graphics* 32, 6 (Nov. 2013), Article No. 211. <https://doi.org/10.1145/2508363.2508422>
- Luis E. Gamboa, Jean-Philippe Guertin, and Derek Nowrouzezahrai. 2018. Scalable Appearance Filtering for Complex Lighting Effects. *ACM Trans. Graph.* 37, 6, Article 277 (Dec. 2018), 13 pages. <https://doi.org/10.1145/3272127.3275058>
- Johannes Hanika, Marc Droske, and Luca Fascione. 2015a. Manifold Next Event Estimation. *Comput. Graph. Forum* 34, 4 (July 2015), 87–97. <http://dl.acm.org/citation.cfm?id=2858834.2858844>
- Johannes Hanika, Anton Kaplanyan, and Carsten Dachsbacher. 2015b. Improved Half Vector Space Light Transport. *Comput. Graph. Forum* 34, 4 (July 2015), 65–74. <http://dl.acm.org/citation.cfm?id=2858834.2858842>
- Eric Heitz. 2014. Understanding the Masking-Shadowing Function in Microfacet-Based BRDFs. *Journal of Computer Graphics Techniques (JCGT)* 3, 2 (30 June 2014), 48–107. <http://jcggt.org/published/0003/02/03/>
- Eric Heitz. 2017. *Geometric Derivation of the Irradiance of Polygonal Lights*. Research Report. Unity Technologies. <https://hal.archives-ouvertes.fr/hal-01458129>

- Eric Heitz, Jonathan Dupuy, Stephen Hill, and David Neubelt. 2016. Real-Time Polygonal-Light Shading with Linearly Transformed Cosines. *ACM Trans. Graph.* 35, 4, Article 41 (July 2016), 8 pages. <https://doi.org/10.1145/2897824.2925895>
- Sebastian Herholz, Oskar Elek, Jiří Vorba, Hendrik Lensch, and Jaroslav Krivánek. 2016. Product Importance Sampling for Light Transport Path Guiding. *Computer Graphics Forum* 35, 4 (2016), 67–77. <https://doi.org/10.1111/cgf.12950>
- Homan Igehy. 1999. Tracing Ray Differentials. In *Proceedings of the 26th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '99)*. ACM Press/Addison-Wesley Publishing Co., USA, 179–186. <https://doi.org/10.1145/311535.311555>
- Wenzel Jakob. 2013. *Light transport on path-space manifolds*. Ph.D. Dissertation. Cornell University.
- Wenzel Jakob, Miloš Hašan, Ling-Qi Yan, Jason Lawrence, Ravi Ramamoorthi, and Steve Marschner. 2014. Discrete Stochastic Microfacet Models. *ACM Trans. Graph.* 33, 4, Article 115 (July 2014), 10 pages. <https://doi.org/10.1145/2601097.2601186>
- Wenzel Jakob and Steve Marschner. 2012. Manifold Exploration: A Markov Chain Monte Carlo Technique for Rendering Scenes with Difficult Specular Transport. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 31, 4 (July 2012), 58:1–58:13. <https://doi.org/10.1145/2185520.2185554>
- Anton S. Kaplanyan, Johannes Hanika, and Carsten Dachsbacher. 2014. The Natural-constraint Representation of the Path Space for Efficient Light Transport Simulation. *ACM Trans. Graph.* 33, 4, Article 102 (July 2014), 13 pages. <https://doi.org/10.1145/2601097.2601108>
- A. S. Kaplanyan, S. Hill, A. Patney, and A. Lefohn. 2016. Filtering Distributions of Normals for Shading Antialiasing. In *Proceedings of High Performance Graphics (Dublin, Ireland) (HPG '16)*. Eurographics Association, Goslar Germany, Germany, 151–162. <https://doi.org/10.2312/hpg.20161201>
- Alexandr Kuznetsov, Miloš Hašan, Zexiang Xu, Ling-Qi Yan, Bruce Walter, Nima Khademi Kalantari, Steve Marschner, and Ravi Ramamoorthi. 2019. Learning Generative Models for Rendering Specular Microgeometry. *ACM Trans. Graph.* 38, 6, Article 225 (Nov. 2019), 14 pages. <https://doi.org/10.1145/3355089.3356525>
- Eric P. Lafortune and Yves D. Willems. 1993. Bi-directional path tracing. In *Proceedings of Third International Conference on Computational Graphics and Visualization Techniques (Compugraphics '93)*. Alvor, Portugal, 145–153.
- Johann Heinrich Lambert. 1760. *Photometria sive de mensura et gradibus luminis, colorum et umbrae*.
- Don Mitchell and Pat Hanrahan. 1992. Illumination from Curved Reflectors. *SIGGRAPH Comput. Graph.* 26, 2 (July 1992), 283–291. <https://doi.org/10.1145/142920.134082>
- Thomas Müller. 2019. “Practical Path Guiding” in Production. In *ACM SIGGRAPH Courses: Path Guiding in Production, Chapter 10* (Los Angeles, California). ACM, New York, NY, USA, 18:35–18:48. <https://doi.org/10.1145/3305366.3328091>
- Thomas Müller, Markus Gross, and Jan Novák. 2017. Practical Path Guiding for Efficient Light-Transport Simulation. *Computer Graphics Forum* 36, 4 (June 2017), 91–100. <https://doi.org/10.1111/cgf.13227>
- Merlin Nimier-David, Delio Vicini, Tizian Zeltner, and Wenzel Jakob. 2019. Mitsuba 2: A Retargetable Forward and Inverse Renderer. *Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 38, 6 (Dec. 2019). <https://doi.org/10.1145/3355089.3356498>
- Tomoyuki Nishita and Eiichi Nakamae. 1985. Continuous Tone Representation of Three-Dimensional Objects Taking Account of Shadows and Interreflection. In *Proceedings of the 12th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '85)*. Association for Computing Machinery, New York, NY, USA, 23–30. <https://doi.org/10.1145/325334.325169>
- Marc Olano and Dan Baker. 2010. LEAN Mapping. In *Proceedings of the 2010 ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games (Washington, D.C.) (I3D '10)*. ACM, New York, NY, USA, 181–188. <https://doi.org/10.1145/1730804.1730834>
- Christoph Peters and Carsten Dachsbacher. 2019. Sampling Projected Spherical Caps in Real Time. *Proc. ACM Comput. Graph. Interact. Tech.* 2, 1, Article 1 (June 2019), 16 pages. <https://doi.org/10.1145/3320282>
- Matt Pharr, Wenzel Jakob, and Greg Humphreys. 2016. *Physically Based Rendering: From Theory to Implementation (3rd ed.)* (3rd ed.). Morgan Kaufmann Publishers Inc., San Francisco, CA, USA. 1266 pages.
- Boris Raymond, Gaël Guennebaud, and Pascal Barla. 2016. Multi-scale Rendering of Scratched Materials Using a Structured SV-BRDF Model. *ACM Trans. Graph.* 35, 4, Article 57 (July 2016), 11 pages. <https://doi.org/10.1145/2897824.2925945>
- Florian Reibold, Johannes Hanika, Alisa Jung, and Carsten Dachsbacher. 2018. Selective Guided Sampling with Complete Light Transport Paths. *ACM Trans. Graph.* 37, 6, Article 223 (Dec. 2018), 14 pages. <https://doi.org/10.1145/3272127.3275030>
- Vincent Schüssler, Eric Heitz, Johannes Hanika, and Carsten Dachsbacher. 2017. Microfacet-based Normal Mapping for Robust Monte Carlo Path Tracing. *ACM Trans. Graph.* 36, 6, Article 205 (Nov. 2017), 12 pages. <https://doi.org/10.1145/3130800.3130806>
- B. Smith. 1967. Geometrical shadowing of a random rough surface. *IEEE Transactions on Antennas and Propagation* 15, 5 (Sep. 1967), 668–671. <https://doi.org/10.1109/TAP.1967.1138991>
- T. S. Trowbridge and K. P. Reitz. 1975. Average irregularity representation of a rough surface for ray reflection. *J. Opt. Soc. Am.* 65, 5 (May 1975), 531–536. <https://doi.org/10.1364/JOSA.65.000531>
- C. Ureña. 2000. Computation of Irradiance from Triangles by Adaptive Sampling. *Computer Graphics Forum* 19, 2 (2000), 165–171. <https://doi.org/10.1111/1467-8659.00452> arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1111/1467-8659.00452
- Carlos Ureña and Iliyan Georgiev. 2018. Stratified Sampling of Projected Spherical Caps. *Comput. Graph. Forum* 37 (2018), 13–20.
- Eric Veach and Leonidas J. Guibas. 1994. Bidirectional Estimators for Light Transport. In *Eurographics Rendering Workshop 1994 Proceedings*. Darmstadt, Germany, 147–162.
- Eric Veach and Leonidas J. Guibas. 1997. Metropolis Light Transport. In *Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '97)*. ACM Press/Addison-Wesley Publishing Co., New York, NY, USA, 65–76. <https://doi.org/10.1145/258734.258775>
- Jiří Vorba, Ondřej Karlik, Martin Šik, Tobias Ritschel, and Jaroslav Krivánek. 2014. On-line Learning of Parametric Mixture Models for Light Transport Simulation. *ACM Transactions on Graphics (Proceedings of SIGGRAPH 2014)* 33, 4 (aug 2014).
- Bruce Walter, Zhao Dong, Steve Marschner, and Donald P. Greenberg. 2015. The Ellipsoid Normal Distribution Function, supplemental material of Predicting Appearance from Measured Microgeometry of Metal Surfaces. *ACM Trans. Graph.* 35, 1 (2015), 9:1–9:13.
- Bruce Walter, Stephen R. Marschner, Hongsong Li, and Kenneth E. Torrance. 2007. Microfacet Models for Refraction Through Rough Surfaces. In *Proceedings of the 18th Eurographics Conference on Rendering Techniques (Grenoble, France) (EGSR'07)*. Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, 195–206. <https://doi.org/10.2312/EGWR/EGSR07/195-206>
- Bruce Walter, Shuang Zhao, Nicolas Holzschuch, and Kavita Bala. 2009. Single Scattering in Refractive Media with Triangle Mesh Boundaries. *ACM Transactions on Graphics* 28, 3 (Aug. 2009), 92:1–8. <https://doi.org/10.1145/1531326.1531398>
- Sebastian Werner, Zdravko Velinov, Wenzel Jakob, and Matthias B. Hullin. 2017. Scratch Iridescence: Wave-Optical Rendering of Diffractive Surface Structure. *ACM Trans. Graph.* 36, 6, Article 207 (Nov. 2017), 14 pages. <https://doi.org/10.1145/3130800.3130840>
- Chao Xu, Rui Wang, Shuang Zhao, and Hujun Bao. 2017. Real-Time Linear BRDF MIP-Mapping. *Computer Graphics Forum* 36, 4 (2017), 27–34. <https://doi.org/10.1111/cgf.13221> arXiv:https://onlinelibrary.wiley.com/doi/pdf/10.1111/cgf.13221
- Ling-Qi Yan, Miloš Hašan, Wenzel Jakob, Jason Lawrence, Steve Marschner, and Ravi Ramamoorthi. 2014. Rendering Glints on High-resolution Normal-mapped Specular Surfaces. *ACM Trans. Graph.* 33, 4, Article 116 (July 2014), 9 pages. <https://doi.org/10.1145/2601097.2601155>
- Ling-Qi Yan, Miloš Hašan, Steve Marschner, and Ravi Ramamoorthi. 2016. Position-normal Distributions for Efficient Rendering of Specular Microstructure. *ACM Trans. Graph.* 35, 4, Article 56 (July 2016), 9 pages. <https://doi.org/10.1145/2897824.2925915>
- Ling-Qi Yan, Miloš Hašan, Bruce Walter, Steve Marschner, and Ravi Ramamoorthi. 2018. Rendering Specular Microgeometry with Wave Optics. *ACM Trans. Graph.* 37, 4, Article 75 (July 2018), 10 pages. <https://doi.org/10.1145/3197517.3201351>
- Tizian Zeltner, Iliyan Georgiev, and Wenzel Jakob. 2020. Specular Manifold Sampling for Rendering High-Frequency Caustics and Glints. *Transactions on Graphics (Proceedings of SIGGRAPH)* 39, 4 (July 2020). <https://doi.org/10.1145/3386569.3392408>
- Tobias Zirr and Anton S. Kaplanyan. 2016. Real-Time Rendering of Procedural Multi-scale Materials. In *Proceedings of the 20th ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games (Redmond, Washington) (I3D '16)*. Association for Computing Machinery, New York, NY, USA, 139–148. <https://doi.org/10.1145/2856400.2856409>