Eurographics Symposium on Rendering (DL-only Track) (2019)
T. Boubekeur and P. Sen (Editors)

Impulse Responses for Precomputing Light from Volumetric Media

Adrien Dubouchet[1,2], Peter-Pike Sloan[2], Wojciech Jarosz[3], Derek Nowrouzezahrai[4]

[1]Université de Montréal, Canada, [2]Activision Publishing Inc., United States
[3]Dartmouth College, United States, [4]McGill University, Canada

Figure 1: *In-game screenshot (a) with complex surface and volume shading. Standard lightmaps (d) ignore volume-to-surface transport, while directly-attenuated volume-to-surface illumination (b) captures media shadowing only. Our method (c) additionally capture volumetric single- and multiple-scattered light from emitters and surfaces (e). ©Activision Publishing, Inc.*

Abstract

Modern interactive rendering can rely heavily on precomputed static lighting on surfaces and in volumes. Scattering from volumetric media can be similarly treated using precomputation, but transport from volumes onto surfaces is typically ignored here. We propose a compact, efficient method to simulate volume-to-surface transport during lighting precomputation . We leverage a novel model of the spherical impulse response of light scattered (and attenuated) in volumetric media to simulate light transport from volumes onto surfaces with simple precomputed lookup tables. These tables model the impulse response as a function of distance and angle to the light and surfaces. We then remap the impulse responses to media with arbitrary, potentially heterogeneous scattering parameters and various phase functions. Moreover, we can compose our impulse response model to treat multiple scattering events in the volume (arriving at surfaces). We apply our method to precomputed volume-to-surface light transport in complex scenes, generating results indistinguishable from ground truth simulations. Our tables allow us to precompute volume-to-surface transport orders of magnitude faster than even an optimized path tracing-based solution would.

CCS Concepts

• ***Computing methodologies*** → ***Ray tracing;***

1. Introduction

Volumetric participating media can significantly impact the realism of virtual scenes, due to both the subtle interplay between surface- and volume-transport effects and the presence of entities like smoke and clouds. Modern video games increasingly seek to include such effects to improve realism and artistic flexibility. While many works propose efficient simulation methods for volumetric single- and multiple-scattering (e.g., [Hil15; Wro14]), no interactive methods accurately model the effects of scattering *from* a volume *onto* other surfaces/locations in a scene, let alone several such bounces. *Volume-to-surface* transport has been studied in offline image synthesis but these methods are too costly for interactive content pipelines.

Some interactive applications rely heavily on precomputed lighting [Che08; IS17b; McT04; ODo18] where the tradeoffs lie between what to precompute and what to evaluate at run-time. Indirect light and lighting from complex area/sky sources is more often precomputed, whereas run-time light source influence radii are often restricted to reduce shading; here, their artificially-clamped shading may also be precomputed [Cha18], leading to shading that can be interpolated between run-time evaluated and precomputed sources. Precomputed lighting data is cached on surfaces and in free-space: lightmaps, per-vertex data and light probes. Precomputed light probes are used to relight dynamic objects with the effects of the (static) lighting in the scene. Relying on even the most efficient offline methods in order to include the effects of volumetric scattering in these precomputed datasets would require hours to days to converge for even moderately-sized in-game assets.

As such, even in scenes with participating media, the majority of games only precompute surface-to-surface transport and will continue to rely on such methods for the foreseeable future, a trend only exacerbated for mobile and virtual reality platforms due to power constraints and frame rate requirements. We propose a very efficient method to include the neglected volume-to-surface interactions during precomputation, resulting in more accurate lighting with only modest performance overhead. Instead of relying on costly volumetric light transport simulation, we model the *spherical impulse*

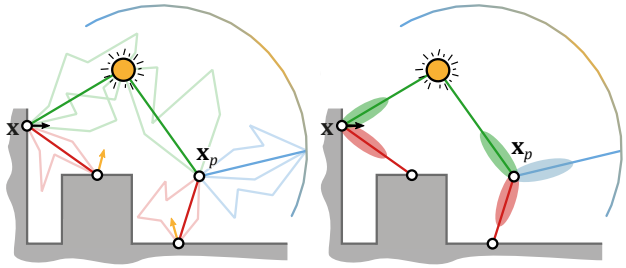

Figure 2: *Consider surface* $\mathbf{x}$ *and probe* $\mathbf{x}_p$ *locations lit by point (green), environment (blue) and indirect (red) light. Offline rendering traces light paths to solve the radiative transport equation (left; transparent lines). We encode the impulse response of volumetric transport (right), to model media scattering in precomputed lighting.*

response of volumetric light in a form that allows us to quickly precompute volume-to-surface transport. Though our method is devised to accelerate bake times, its reliance on per-ray table lookups for approximate inscattered light makes it an attractive and inexpensive alternative to the full series expansion for offline path tracing algorithms as well. Specifically, our contributions are:

- a formulation and analysis of the impulse response of volumetric in-scattered light for point, directional and differential emitters,
- a modular extension from single- to multiple-scattering regimes,
- a treatment of *heterogeneous*, *anisotropic* volume scattering, and
- a compact, pretabulated *zonal harmonic* parameterization of the impulse response applied to a constant-time algorithm for adding accurate volume-to-surface transport to precomputed light maps.

2. Related Work

Interactive rendering and games rely on faithful simulations of physically based light transport effects. While many effects remain outside the performance constraints imposed in these applications, precomputation-based solutions provide an avenue to include complex lighting with low runtime cost. Building on light mapping [Abr00] where precomputed lighting is cached at texels, separate from higher frequency albedo, recent methods cache other forms of light transport data: e.g., by representing variations due to normal mapped surfaces [McT04] in a spherical/hemispherical basis, non-Lambertian view-dependent reflection can also be cached [Che08; IS17a; NP15; ODo18]. This added angular resolution can, in turn, allow baked data to be represented at coarser spatial resolutions, such as per-vertex [KBS11] or per-probe [IS17b], improving storage and performance. For semi-static geometry, lighting is also precomputed in volumetric data structures, variants of irradiance volumes or discrete ordinate methods [CNS*11; GSHG98; KD10]. Precomputation time can take between minutes and hours, and even with new GPU advances it is costly [Hil18].

Precomputed Radiance Transfer (PRT) cache transport operators (e.g., per vertex) to model how incoming light (typically represented in a basis space) is locally transformed by occlusion and scattering. The basis representation allows for precomputed transport to be adapted, at runtime, to dynamic lighting and/or viewing scenarios [Leh07; LWDB10; SHHS03; SKS02; ZDM13]. Several approaches explore various discretizations of the caching domain, including methods that optimize the placement and interpolation of per-vertex lighting data [KBS11] and lighting "probes" placed in the open space of a scene in order to relight dynamic objects [GSHG98; SL17]. These precomputation-based methods allow for costly effects, like interreflections or ambient occlusion, to be incorporated at a fraction of the cost required for their dynamic evaluation.

Many basis choices have been used for PRT. Spherical and Zonal Harmonics (SH, ZH) benefit from closed-form rotation and convolution operations [RH01; SKS02; Slo08; SLS05]. They are unable, however, to compactly represent higher frequency angular variation. On the other hand, wavelet bases are capable of representing angular variation across frequencies, but at the cost of more complicated runtime operations [HPB06]. Volumetric transport tends towards angularly smooth radiance distributions, and so our work relies on compact ZH formulations to encode the impulse response of emitted light scattered within a volumetric media.

In volumetric shading, the underlying transport theory and effective Monte Carlo-based solutions are well understood [dEon16; KF12; NGHJ18; SRNN05], and a tremendous amount of work has found success in applications to offline visual effects [FWKH17; WB11]. Path tracing has dominated film rendering [CJ16; FWKH17; PJH16] due to its ability to scale and its speed to initial (albeit noisy) image pixels. Two dominant forms of caching have furthered the adoption of these methods: radiance caching [JDZJ08; JZJ08a; MJJG18; SJJ12] and photon mapping [BJ17; DJBJ19; JC98; JNSJ11; JNT*11; JZJ08b]. In contrast, the performance constraints of interactive graphics necessitate the use of simplified volume shading methods [ERDS14; Wro14].

Early work on subsurface scattering [JMLH01] modeled variation of scattering parameters by assuming local homogeneity. Our method approximate heterogeneous media similarly, relying on scattering parameters along a ray to make a *directional-homogeneity* simplification. Tabulated BSSRDFs, and their applications to discrete random media [MWM07] and shell transport functions [MPG*16] are also tangentially related to our impulse response models for volume-to-point transport.

3. Background and Theoretical Model

We review volumetric light transport (Section 3.1) and present a factorization we will rely on (Section 3.2), before providing important definitions and properties of the SH basis (Section 3.3). We focus on efficient simulation of costly volume-to-point transport effects during surface lightmap and lightprobe precomputation.

3.1. Light Transport in Volumetric Media

Light transport in volumetric participating media is governed by the *radiative transfer equation* (RTE) [Cha60], a differential equation that describes the change in radiance L along a ray due to scattering and absorption. Integrating this equation and treating the surface *rendering equation* [ICG86; Kaj86] as a boundary condition yields the *volume rendering equation* (VRE) for the radiance $L(\mathbf{x}, \boldsymbol{\omega})$ arriving

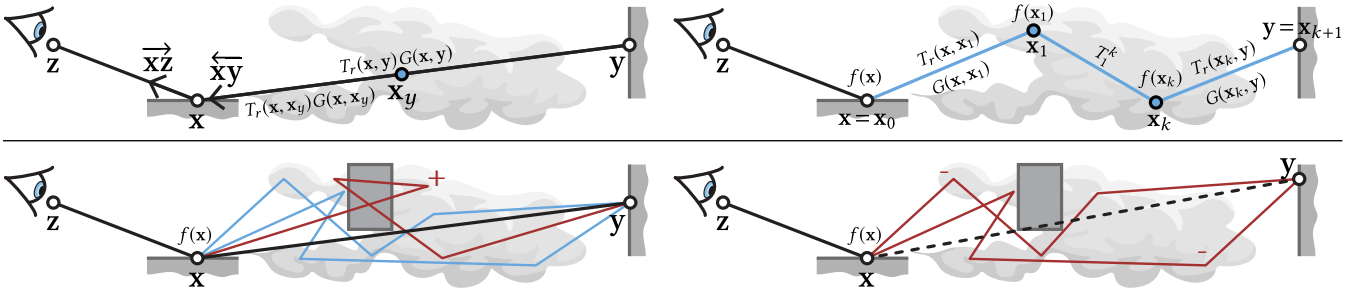

Figure 3: *Top left: 3-point geometry for radiance reflecting at* $\mathbf{x}$ *towards* $\mathbf{z}$ *as the sum of radiance arriving from surface point* $\mathbf{y}$ *and all media points* $\mathbf{x}_y$*. Top right: expanding the recursion results in path throughput of a media subpath between* $\mathbf{x}$ *and* $\mathbf{y}$*. Bottom left: we decompose the incident radiance at* $\mathbf{x}$ *as the sum of the attenuated surface radiance from* $\mathbf{y}$ ***(black)*** *and radiance arriving at* $\mathbf{x}$ *indirectly from* $\mathbf{y}$ *with any number of media bounces in between (blue). This can include false contributions from light paths which should be blocked (red +). Bottom right: We ignore contributions from surfaces that are directly occluded, even though radiance could still arrive at* $\mathbf{x}$ *indirectly from* $\mathbf{y}$ *(red −).*

at point $\mathbf{x}$ from direction ω:

$$L(\mathbf{x},\omega) = \underbrace{\int_0^s \sigma_t(\mathbf{x}_t)\, T_r(\mathbf{x},\mathbf{x}_t)\, L_o(\mathbf{x}_t,\omega)\, \mathrm{d}t}_{L_m(\mathbf{x},\omega)} + \underbrace{T_r(\mathbf{x},\mathbf{x}_s)\, L_o(\mathbf{x}_s,\omega)}_{L_s(\mathbf{x},\omega)}, \quad (1)$$

where L_m is the radiance arriving from all points $\mathbf{x}_t = \mathbf{x} - t\omega$ in the medium, L_s is the radiance arriving from a surface at $\mathbf{x}_s = \mathbf{x} - s\omega$, and σ_t is the media's extinction coefficient. Transmittance T_r models the attenuation due to extinction between two points:

$$T_r(\mathbf{x},\mathbf{x}_t) = \mathrm{e}^{-\int_0^t \sigma_t(\mathbf{x}+t'\omega)\,\mathrm{d}t'} = \mathrm{e}^{-\sigma_t t}. \quad (2)$$

The outgoing radiance L_o is the sum of emitted radiance L_e and the angular integral of the incident radiance:

$$L_o(\mathbf{x},\omega) = L_e(\mathbf{x},\omega) + \int_{\mathcal{S}^2} f(\mathbf{x},\omega,\omega_i)\, L(\mathbf{x},\omega_i)\, \mathrm{d}_\perp\omega_i, \quad (3)$$

where ω_i and ω are incident and outgoing spherical directions about $\mathbf{x}$, $\mathrm{d}_\perp\omega_i$ is the differential (projected) solid angle, and f is either the volumetric phase function $f_p(\mathbf{x},\omega,\omega_i)$ or surface BSDF $f_r(\mathbf{x},\omega,\omega_i)$, depending on whether $\mathbf{x}$ is in the medium $\mathcal{V}$ or on a surface $\mathcal{A}$:

$$f(\mathbf{x},\omega,\omega_i) = \begin{cases} \alpha(\mathbf{x}) f_p(\mathbf{x},\omega,\omega_i) & \text{if } \mathbf{x}\in\mathcal{V}, \\ f_r(\mathbf{x},\omega,\omega_i) & \text{if } \mathbf{x}\in\mathcal{A}, \end{cases} \quad (4)$$

where $\alpha = \sigma_s/\sigma_t$ is the albedo and σ_s the scattering coefficient. We precompute and tabulate response parameterized by the scattering coefficient σ_s; we devise an approach to optionally incorporate absorption coefficient during precomputation (Section 4.3).

3.2. Radiance Decomposition

We can reparameterize Eq. (3) in a three-point configuration as

$$L_o(\overrightarrow{\mathbf{xz}}) = L_e(\overrightarrow{\mathbf{xz}}) + \int_{\mathcal{A}} f(\mathbf{x})\, L(\overleftarrow{\mathbf{xy}})\, \mathrm{d}A(\mathbf{y}) \quad (5)$$

where we let $\overrightarrow{\mathbf{xz}}$ (and $\overleftarrow{\mathbf{xz}}$) denote the unit direction from $\mathbf{x}$ to $\mathbf{z}$ (and from $\mathbf{z}$ to $\mathbf{x}$), $L(\overleftarrow{\mathbf{xy}}) = L_s(\overleftarrow{\mathbf{xy}}) + L_m(\overleftarrow{\mathbf{xy}})$ is the radiance arriving at $\mathbf{x}$ from surface points $\mathbf{y}\in\mathcal{A}$, with

$$L_s(\overleftarrow{\mathbf{xy}}) = T_r(\mathbf{x},\mathbf{y})\, G(\mathbf{x},\mathbf{y})\, L_o(\overrightarrow{\mathbf{yx}}), \quad \text{and} \quad (6)$$

$$L_m(\overleftarrow{\mathbf{xy}}) = \int_0^y T_r(\mathbf{x},\mathbf{x}_y)\, G(\mathbf{x},\mathbf{x}_y)\, L_o(\overrightarrow{\mathbf{x}_y\mathbf{x}})\, \mathrm{d}y, \quad (7)$$

where $\mathbf{x}_y = \mathbf{x} + y\cdot\overrightarrow{\mathbf{xy}} \in \mathcal{V}$ are points in the volume between $\mathbf{x}$ and $\mathbf{y}$, and G and f are generalized geometry and scattering terms:

$$G(\mathbf{x},\mathbf{y}) = \frac{D_\mathbf{x}(\mathbf{y})\, V(\mathbf{x},\mathbf{y})\, D_\mathbf{y}(\mathbf{x})}{\|\mathbf{x}-\mathbf{y}\|^2},\ D_\mathbf{x}(\mathbf{y}) = \begin{cases} |\vec{n}_\mathbf{x}\cdot\overrightarrow{\mathbf{xy}}| & \text{if } \mathbf{x}\in\mathcal{A}, \\ 1 & \text{if } \mathbf{x}\in\mathcal{V}, \end{cases} \quad (8)$$

$$f(\mathbf{x}_i) = \begin{cases} \sigma_t(\mathbf{x}_i)\,\alpha(\mathbf{x}_i)\, f_p(\mathbf{x}_{i+1},\mathbf{x}_i,\mathbf{x}_{i-1}) & \text{if } \mathbf{x}\in\mathcal{V}, \\ f_r(\mathbf{x}_{i+1},\mathbf{x}_i,\mathbf{x}_{i-1}) & \text{if } \mathbf{x}\in\mathcal{A}. \end{cases} \quad (9)$$

Here, V is binary visibility and $\vec{n}_\mathbf{x}$ the normal at $\mathbf{x}$ (Fig. 3, top left).

Assuming non-emissive media and expanding the recursion at media scattering events in Eq. (7), we can rewrite $L(\overleftarrow{\mathbf{xy}})$ as a sum of the radiance arriving at $\mathbf{x}$ from points $\mathbf{y}$ with ≥ 0 medium scattering events in between:

$$L(\overleftarrow{\mathbf{xy}}) = \sum_{k=0}^{\infty} \underbrace{\int \dots \int T_0^k\, L_o(\overrightarrow{\mathbf{yx}_k})\, \mathrm{d}V(\mathbf{x}_1)\dots\mathrm{d}V(\mathbf{x}_k)}_{L^k(\mathbf{x},\mathbf{y})} \quad (10)$$

where $L^k(\mathbf{x},\mathbf{y})$ is incident radiance at $\mathbf{x}$ that has scattered **exactly** k **times** in the medium after leaving $\mathbf{y}$, and T_0^k is the throughput for subpath $\mathbf{x}_0 \equiv \mathbf{x},\dots,\mathbf{x}_{k+1} \equiv \mathbf{y}$ (see Fig. 3; top right, bottom left):

$$T_0^k = \left[\prod_{i=1}^{k} f(\mathbf{x}_i)\right]\left[\prod_{i=0}^{k} T_r(\mathbf{x}_i,\mathbf{x}_{i+1})\, G(\mathbf{x}_i,\mathbf{x}_{i+1})\right]. \quad (11)$$

Note, for $k=0$, Eq. (10) reduces to the attenuated surface radiance in Eq. (6), whereas $k>0$ accounts for recursively expanded media radiance in Eq. (7). Moreover, while $L^k(\mathbf{x},\mathbf{y})$ represents radiance arriving (potentially indirectly) from $\mathbf{y}$, the incident direction at $\mathbf{x}$ is defined by $\overleftarrow{\mathbf{xx}_1}$, which can be any spherical direction about $\mathbf{x}$ when $k\neq 0$. This means that when Eq. (10) is inserted into Eq. (5), the scattering function should be evaluated as $f(\mathbf{x}_1,\mathbf{x},\mathbf{z})$. When $k=0$, $\mathbf{x}_1 = \mathbf{x}_{k+1} = \mathbf{y}$, so the scattering is evaluated as $f(\mathbf{y},\mathbf{x},\mathbf{z})$.

3.3. Spherical Harmonics – Definitions and Properties

Preliminaries. Consider a scalar-valued function $f(\omega)$ over the unit sphere $\mathcal{S}^2$, with spherical directions $\omega = (x,y,z) = (\theta,\phi) \in S^2$. We obtain a vector $\boldsymbol{f}$ of SH coefficients by projecting $f(\omega)$ onto the real SH basis as $\boldsymbol{f} = \int_{S^2} f(\omega)\, \boldsymbol{y}(\omega)\mathrm{d}\omega$, where $\boldsymbol{f} = \{f_0^0, f_1^{-1}, f_1^0, f_1^1, \dots\}$ is a vector of scalar projection coefficients f_l^m and $\boldsymbol{y}(\omega) =$

$\{y_0^0, y_1^{-1}, y_1^0, y_1^1, \ldots\}$ a vector of the SH basis functions:

$$y_l^m(\theta,\phi) = \begin{cases} \sqrt{2}\, K_l^m \, \sin(|m|\phi)\, P_l^{|m|}(\cos\theta), & m \le 0 \\ \sqrt{2}\, K_l^m \, \cos(m\phi)\, P_l^m(\cos\theta) & , \; m > 0 \end{cases}, \qquad (12)$$

where P_l^m are associated Legendre polynomials and K_l^m are normalization factors. An order-N expansion of f onto the SH basis includes all functions for bands $0 \le l \le N-1$. Each band comprises $2l+1$ basis functions, indexed by m. For a fixed band l, each of the basis functions is a degree l polynomials in the Cartesian coordinates (x,y,z) of the (unit) direction ω, and we often rely on a single indexing scheme, with $i = l(l+1)+m$, for brevity and convenience.

Zonal Harmonics. Zonal $m=0$ subset of SH functions, $y_l^0(\omega) = y_l^0(\theta)$, are circularly symmetric about $\cos\theta = z$ and referred to as *zonal harmonics* (ZH). SLOAN et al. [SLS05] introduced a fast rotation convolutional formulation for ZHs in order to align any weighted combination of the y_l^0 functions about an *arbitrary* axis $\bar{\omega} \neq x$. Doing so yield a function that can no longer be reconstructed as using solely a weighted combination of ZH functions, however the SH coefficients h_l^m of this arbitrarily-aligned (circularly symmetric) function can be directly obtained from the *ZH projection coefficients* f_l of the function in its original orientation (i.e.,about z), as

$$h_l^m = n_l^* \, f_l \, y_l^m(\bar{\omega}) = f_l^* \, y_l^m(\bar{\omega}) \,, \qquad (13)$$

where $n_l^* = \sqrt{4\pi/(2l+1)}$ are convolutional normalization factors arising from the fact that Eq. (13) is the Funke-Hecke theorem applied to the original zonal function and a delta at $\bar{\omega}$.

Double-product Integration. Given two spherical functions $a(\omega)$ and $b(\omega)$ with (effective) bandlimits N_a and N_b, the integral of their product is $\int_{S^2} a(\omega) b(\omega) d\omega = \sum_{i=0}^{n^2-1} a_i \, b_i$, where $n = \min(N_a, N_b)$ and we arrive at the RHS by substituting the SH expansions of a and b into the LHS and then applying the orthonormality property of SH: $\int_{S^2} y_i(\omega) y_k(\omega) \mathrm{d}\omega = \delta_{i,k}$, where $\delta_{i,k}$ is the Kroenecker delta.

We derive a double-product formulation for the special-case where a and b are (arbitrarily rotated) circularly symmetric zonal functions (see Appendix A), and we leverage this more efficient integration formulation in our shading algorithm (Section 4).

4. Impulse Response Models & Extensions

We present our theoretical contributions below, leaving discussions of practical and technical considerations for Section 5. Specifically,

- we factorize homogeneous inscattered radiance (Section 4.1) to motivate a novel *spherical impulse response* formulation for volume-to-point transport with three emission schemes (point-, surface- and indirectly-reflected cosine-profiles; Section 4.2),
- we extend these canonical impulse responses to support media with heterogeneous scattering coefficients (Section 4.4) and anisotropic phase functions (Section 4.5), and
- we derive an efficient double-product integration scheme (Appendix A) to decouple incident and outgoing radiance, leading to important run-time flexibility (Section 5).

We aim to accelerate the precomputation of volume-to-point transport in the digital content creation pipeline for interactive graphics applications. We leverage a novel factorized formulation that admits an efficient basis-space integration scheme. Our scheme models singly- and multiply-scattered volumetric transport effects, as well as the effects of extinction through the media (Section 4.1). We avoid costly numerical path-based Monte Carlo simulations typical to existing pipelines used, e.g., in feature film production.

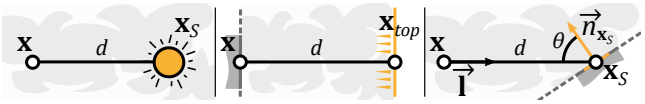

Figure 4: *The three canonical emission types we support: point, directional and differential surface area profiles.*

We precompute small impulse response tables once in a scene-agnostic manner (Sections 4.2 to 4.5), and parameterize these tables to permit fast queries during *surface transport-only* precomputation to incorporate volume-to-point transport effects, and all at a minimal performance overhead[~4-15%](Section 5). Our results closely track ground truth obtained with volumetric path-tracing, despite operating several orders of magnitudes faster (Section 6). Indeed, existing volume-to-point transport simulation incurs a cost so large that it is normally omitted from most interactive graphics pipelines due to the impact on art direction and design iteration times.

4.1. Caching Lightmap Radiance in the Presence of Media

Consider the spherical radiance distribution $L(\mathbf{x},\omega)$ *incident* at location $\mathbf{x} \in \{\mathcal{V} \cup \mathcal{A}\}$ either a lightmap texel or a location in free-space.

Instead of accumulating the contribution of many paths scattering at surfaces and in the volume, and this for many cached receiver locations $\mathbf{x}$, we will consider a bounce- and emitter-based decomposition of the problem. Specifically, by carefully parameterizing the *single-scattered* volumetric radiance arriving at a point from different emitter and reflector configurations, we will first show:

- how to compactly express the incident radiance distribution $L(\mathbf{x},\omega)$ for an *arbitrary* receiver-emitter configuration, for the three emission profiles (see Fig. 4), and
- how to shade lightmaps with this compact representation, composing the single-scattered response to analytically account for approximate multiple-scattering effects.

We will show that zonal harmonics are an effective basis for representing, shading, and sampling from the incident volume-to-point radiance distributions at cache locations in Section 4.2.

4.2. Spherical Impulse Response for Airlight Integrals

Consider the directly attenuated and single-scattered light arriving at $\mathbf{x}$ from a point emitter at $\mathbf{x}_s$ in a homogeneous, isotropic medium. The spherical incident radiance at $\mathbf{x}$ due to in-scattering and extinction $L^{(\le 1)}(\mathbf{x},\omega)$, i.e., from light that has scattered zero and one times before arriving at $\mathbf{x}$, is a special case of Eqs. (1) and (3) with $L_o = L_e$ where the only source radiance is from emitters, $L_e(\mathbf{x},\omega) \neq 0$ iff $\mathbf{x} = \mathbf{x}_s$.

In the case of point emitters, this spherical incident radiance exhibits circular symmetry about the axis from $\mathbf{x}$ to $\mathbf{x}_S$, and so we

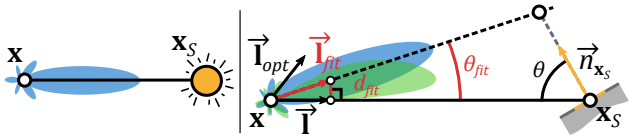

Figure 5: *Light scatterd from point (and directional) sources result in circularly symmetric radiance distributions at* $\mathbf{x}$ *(left; in blue). For differential area sources the distribution skews with the emitters normal and is no-longer circularly symmetric; we determine a best-fit circularly symmetric approximation (right; blue, from green).*

can represent it as $L^{(\leq 1)}(\mathbf{x},\omega) \approx \sum_l L_l^{(\leq 1)} y_l^0(\theta)$ with ZH coefficients

$$L_l^{(\leq 1)} = \frac{\sigma_t}{4\pi}\int_{\mathcal{S}^2}\left[\int_0^s T_r(\mathbf{x},\mathbf{x}_t)\,T_r(\mathbf{x}_t,\mathbf{x}_s)L_e(\mathbf{x}_s,\omega)\,dt + T_r(\mathbf{x},\mathbf{x}_s)\,L_e(\mathbf{x}_s,\omega)\right] y_l^0(\omega)\,d\omega, \quad (14)$$

where $L_e(\mathbf{x}_s,\omega) = I_0/(|\mathbf{x}-\mathbf{x}_S|^2)$ and I_0 is the point source's intensity. For a fixed receiver $\mathbf{x}$, emitter $\mathbf{x}_s$ and media $\{\sigma_t,\sigma_s\}$ configuration we compute the coefficients with Monte Carlo integration, importance sampling points $\mathbf{x}_t$ on the ray in the inner integrand with equi-angular sampling [KF12], and uniformly sample the outer spherical integrand with low-discrepancy quasi-Monte Carlo samples [PJH16]. More advanced strategies exist, taking advantage of special parameterizations of SH [JCJ09] or semi-analytic integration schemes [BXH*18], but we found this yields acceptable results.

We can similarly express the ZH-projected incident radiance at $\mathbf{x}$ due to directly attenuated and single-scattered light arriving from a *directional emitter*. It is important to note that the effects of transmission due to such an infinitely-distant source depends on the projection of the emitter's direction onto the boundary of the medium's bounding shape, which we denote $\mathbf{x}_{top}$ (Fig. 4, middle). For infinite media, the transmittance from *any point* above $\mathbf{x}_{top}$ to the "location" of the directional source is 0. The projection coefficients are,

$$L_l^{(\leq 1)} = \frac{\sigma_t}{4\pi}\int_{\mathcal{S}^2}\left[\int_0^s T_r(\mathbf{x},\mathbf{x}_t)\,T_r(\mathbf{x}_t,\mathbf{x}_{top})L_e(\mathbf{x}_{top},\omega)\,dt + T_r(\mathbf{x},\mathbf{x}_{top})\,L_e(\mathbf{x}_{top},\omega)\right] y_l^0(\omega)\,d\omega \quad (15)$$

and we can compute Eq. (15) with the same MC scheme as before, with $L_e(\mathbf{x}_{top},\omega) = I_0\,\delta(\omega - \overrightarrow{\mathbf{x}\mathbf{x}_{top}})$ for distant directional sources.

The last emitter profile we consider is a differential area source centered at $\mathbf{x}_s$ with normal $\vec{n}_{\mathbf{x}_s}$ (Fig. 4, right). Unlike the point and directional emitters, we will use this source to not only model the effects of area lights but also the effects due to indirectly-reflected surface reflection from global illumination. Furthermore, the directly-attenuated and single-scattered incident radiance at $\mathbf{x}$ due to this source is not generally a circularly-symmetric function: as the source's normal deviates from the direction towards the receiver $\overleftarrow{\mathbf{x}\mathbf{x}_s}$, the radiance's spherical anisotropy increases, necessitating a full SH representation as opposed to a much more compact ZH one.

Fortunately, we observe strongly unimodal (albeit off-axis) directionality in this distribution, even as the angle between $\vec{n}_{\mathbf{x}_s}$ and $\overleftarrow{\mathbf{x}\mathbf{x}_s}$ approaches $\pi/2$ (Fig. 5, right), motivating the following circularly-symmetric approximation: beginning from the full SH projection

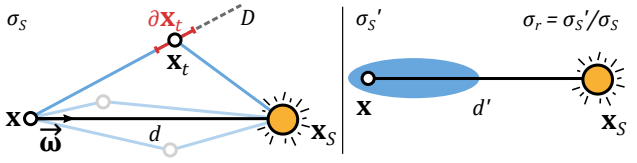

Figure 6: *Left: we compute and store the ZH impulse response at* $\mathbf{x}$ *for many distances* d *and a canonical* σ_s*. Right: when querying the tables in a scene with a different* σ_s'*, we adjust the lookup distance* $d_q = \sigma_r\,d'$ *and scale the ZH coeffs. according to our* σ*-ratio* σ_r*.*

vector of the incident radiance at $\mathbf{x}$,

$$\mathbf{L}^{(\leq 1)} = \frac{\sigma_t}{4\pi}\int_{\mathcal{S}^2}\left[\int_0^s T_r(\mathbf{x},\mathbf{x}_t)\,T_r(\mathbf{x}_t,\mathbf{x}_s)L_e(\mathbf{x}_s,\omega)\,\lfloor\vec{n}_{\mathbf{x}_s}\cdot\omega\rfloor\,dt + T_r(\mathbf{x},\mathbf{x}_s L_e(\mathbf{x}_s,\omega)\,\lfloor\vec{n}_{\mathbf{x}_s}\cdot\omega\rfloor\right]\mathbf{y}(\omega)\,d\omega,$$

one choice for the ZH approximation of $\mathbf{L}^{(\leq 1)}$ uses the *optimal linear direction* $\overrightarrow{\mathbf{l}_{opt}} = (-(L^{(\leq 1)})_1^1, -(L^{(\leq 1)})_1^{-1}, (L^{(\leq 1)})_1^0)$ as the axis of symmetry [SLS05]. We found that the lowest-error ZH approximation's axis depends on the angle θ between $\vec{n}_{\mathbf{x}_s}$ and $\overleftarrow{\mathbf{x}\mathbf{x}_s}$, lying between the direction to the light $\overrightarrow{\mathbf{l}} \equiv -\overleftarrow{\mathbf{x}\mathbf{x}_s}$ and the optimal linear direction $\overrightarrow{\mathbf{l}_{opt}}$. We apply Brent's parabolic interpolation method to search for the best-fit ZH axis $\overrightarrow{\mathbf{l}_{fit}}$, repeating the search for many discretized incident angles. We store the angle θ_{fit} between $\overrightarrow{\mathbf{l}_{fit}}$ and $\overrightarrow{\mathbf{x}\mathbf{x}_s}$ and its ZH coefficients $L_{l,\theta}^{(\leq 1)}$ for every incident angle θ (Fig. 5, right),

$$L_{l,\theta}^{(\leq 1)} = \frac{\sigma_t}{4\pi}\int_{\mathcal{S}^2}\left[\int_0^s T_r(\mathbf{x},\mathbf{x}_t)\,T_r(\mathbf{x}_t,\mathbf{x}_s)\,\lfloor\vec{n}_{\mathbf{x}_s}\cdot\omega\rfloor L_e(\mathbf{x}_s,\omega)\,dt + T_r(\mathbf{x},\mathbf{x}_s)\,L_e(\mathbf{x}_s,\omega)\right] y_l^0\left((\overrightarrow{\mathbf{l}_{fit}}(\theta))\cdot\omega\right)\,d\omega, \quad (16)$$

and $L_e(\mathbf{x}_s,\omega) = I_0/(|\mathbf{x}-\mathbf{x}_S|^2)$. We re-parameterize according to a *shift distance* $d_{fit} = \sin\theta_{fit}$ instead of the angle θ_{fit} to optimize the runtime. These impulse responses only hold for a single receiver-emitter and media configuration. They also only treat homogeneous media with isotropic scattering. We show how to resolve each of these limitations in Sections 4.3 to 4.5 to compactly represent spherical impulse responses for *arbitrary* configurations, before detailing how to apply these responses when computing full, multi-bounce volume-to-point transport for precomputed lighting in Section 5.

4.3. Compact Impulse Response Parameterizations

We reduce from the 6D spatial configuration (3D for each of $\mathbf{x}$ and $\mathbf{x}_s$) to 1D ZH tables for the point and directional sources using the relative distance $d_{\mathbf{x},\mathbf{x}_S} \equiv |\mathbf{x}-\mathbf{x}_S|$ between a receiver/shading point $\mathbf{x}$ and the point emitter location, and the latter by the distance $d_{\mathbf{x},\mathbf{x}_{top}} \equiv |\mathbf{x}-\mathbf{x}_{top}|$ to the nearest position on the boundary of the media $\mathbf{x}_{top}$ towards the directional source. The 2D ZH/d_{fit} tables for differential area sources are parameterized by the relative distance $d_{\mathbf{x},\mathbf{x}_S}$ **and** the (cosine of the) angle θ formed between the source area patch's normal $\vec{n}_{\mathbf{x}_s}$ and the direction $-\overrightarrow{\mathbf{l}}$ from the receiver to the patch (see Fig. 4). While this works for fixed media parameters σ_t and σ_s, 3D and 4D tables would still be unwieldy.

Relative Media Coefficient Parameterization. Next, we eliminate the dependence on media parameter configurations by showing

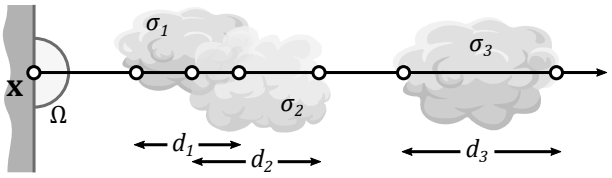

Figure 7: *Heterogeneous media is approximated by spherically-homogeneous media from the point of view of* **x**. *Here, the ray's optical thickness is* $d_1\sigma_1 + d_2\sigma_2 + d_3\sigma_3$. *The ray's scattering coefficient is obtained by normalizing optical thickness by ray length.*

how to analytically map between impulse response coefficients tabulated from a canonical $\{\sigma_t, \sigma_s\}$ to those needed for an arbitrary medium's $\{\sigma'_t, \sigma'_s\}$. To do so, we consider the manner in which solutions to the RTE behave when these parameters change. Begin by considering the impact of changing σ_s. Here, both the transmittance T_r and inscattering probability in Eq. (7) are affected. For $T_r(\mathbf{x}, \mathbf{y}) = T_r(d)$ with $d \equiv \|\mathbf{x} - \mathbf{y}\|$, the optical thickness $-\sigma_s d$ behaves predictably as σ_s changes: modifying the optical thickness for varying σ_s amounts to inversely modifying d. Specifically, $\exp -\sigma_s d = \exp -\sigma'_s d'$ when $d = (\sigma'_s/\sigma_s)d'$. We call $\sigma_r \equiv \sigma'_s/\sigma_s$ the σ-ratio. From this relation, we can precompute the inscattering impulse response for some canonical σ_s and account for changes in the actual media parameters in two simple steps: first, we adjust the table distance query as $d_q = \sigma_r d'$; second, as inscattered events have new scattering probability σ'_s, we account for this discrepancy by scaling the queried ZH coefficients by σ_r.

With point and differential area sources, we need to also consider the effect on emission intensity I_0/d^2, which fall-off with an inverse square profile. When querying the tables with an adjusted d_q, the intensity fall-off must be similarly adjusted by the σ-ratio, as per $I_0/[\sigma_r d]^2$. This adds an additional factor of σ_r^2 for these types of sources that needs to be modeled when computing their inscattered contribution over a single-scattered ray in Eqs. (1) and (3).

When integrating over the inscattered ray and over a distance adjusted by σ_r, a factor of σ_r arises from the integration and cancels one of the above $1/\sigma_r$ intensity adjustment terms. As such, point and differential area entries are ultimately adjusted by a factor of σ_r^2, and directional source entries remain unaffected. Figure 6 summarizes the adjustments necessary for arbitrary media parameters.

Accounting for Absorption. While less often used in game assets, we can also support media absorption and as such, $\sigma_t = \sigma_a + \sigma'_s$ is used. The adjusted query distance and coefficient weighing scheme need only rely on the appropriately modified σ-ratio of $\sigma_r = \sigma_t/\sigma_s$.

We note that including absorption will preclude the efficient table composition method we propose for multiple-scattering in Section 5.1. We will show that, since only the scattering coefficient drives inscattering probability, in the absence of absorption we can arrive at a single table that encodes the ZH response for *all of the media interactions*: direct attenuation, single-scattering and an arbitrary number of multiple-scattering bounces. To do so, we will employ a multi-bounce generalization of the σ-ratio.

Next we discuss how to extend the impulse response formulations to heterogeneous (Section 4.4), anisotropic (Section 4.5) media.

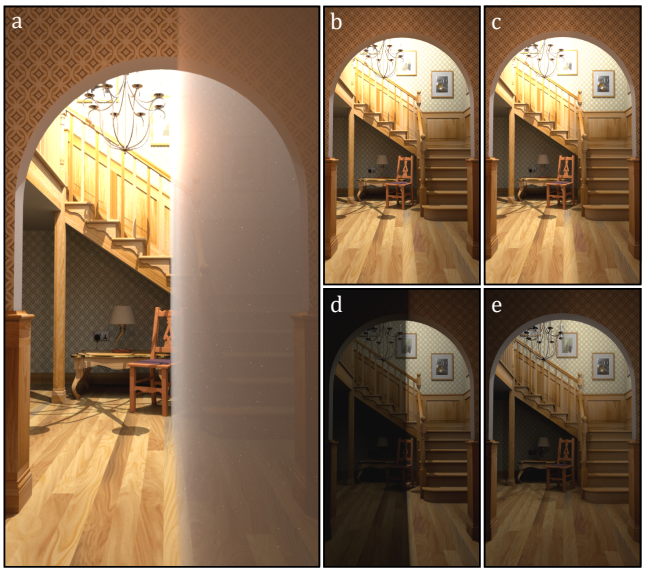

Figure 8: *Comparing our method to ground truth generated in Mitsuba [Jak13] in an extreme scenario. The final rendering (a) includes eye-ray attenuation atop the final precomputed surface lightmaps. Our lightmap (b) and the ground truth lightmap (c) include volume-to-surface transport from directly attenuated (but unoccluded, in our case) lights, and single- and multiple-scattering. Artifacts are most evident when only visualizing inscattering (d) versus ground truth (e): these are due to the extreme heterogeneous discontinuity, the positioning of light sources outside the medium, and surface-to-light connections that do not receive inscattering when not intersecting the volume.*

4.4. Impulse Response in Heterogeneous Media

In games, homogeneous media are most often used to model the effects of atmospheric scattering, but in closed environments the predominant form of participating media used are heterogeneous – with spatially-varying density and scattering parameters. Artists author this media with many strategies, including layering procedural patterns (e.g., Perlin noise), storing and combining volumetric "brushes" with hierarchical data structures (e.g., oriented bounding boxes, octrees), or by "sweeping" 2D image stencils (e.g., gradients, exponential ramps, binary masks).

One of the largest complications due to heterogenous media is the added complexity of evaluating the transmittance T_r, as this necessitates some ray marching-based solution in the setting of general heterogeneity. We make the following simplification to enable another remapping of our (homogeneous) impulse response coefficients to heterogeneous media: we assume that the media acts homogeneously for any fixed ray in space, but that the parameters of the per-ray "effective homogeneity" can vary per ray. This can be interpreted as replacing heterogeneous media with a spatio-directionally-varying homogeneous media: every ray in a scene "observes" a different homogeneous medium. The effective scattering parameter can be efficiently computed: $\sigma_s^{\text{effective}} = (\sum_i \sigma_{s,i}\, d_i)/D$, where D is the total ray length (Fig. 7).

We can now apply σ-ratio remapping to our impulse response tables (Section 4.3) with $\sigma_s^{\text{effective}}$ to query the table entries. Note

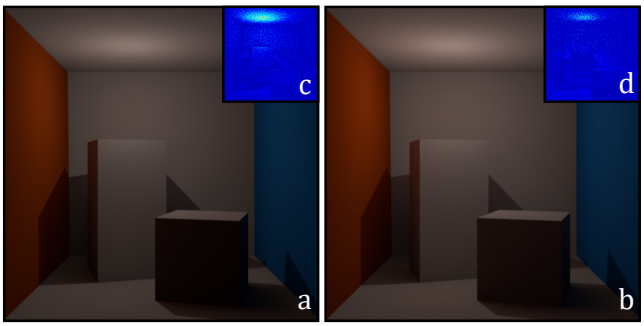

Figure 9: *Cornell box with homogeneous media, 10-bounce inscattering-only without (a) and with (b) back-scattering. (c) and (d) show false-color difference to unconverged PT at 3EV.*

that our assumption for heterogeneous media does not introduce any approximation for the directly-attenuated volumetric light, since the product of transmittance (for discrete changes in density) along the ray is equal to the transmittance of the average density over the ray: $\exp(-\sum_i \sigma_{s,i}\, d_i)/D = \Pi_i \exp(-\sigma_{s,i}\, d_i)$. The single-scattered impulse contribution will, however, incur approximation error since it integrates rays that fall off-axis from the central ray used to estimate the effective homogeneous scattering coefficient. We only noticed visible artifacts in degenerate cases, e.g., binary media; see Fig. 8 for an failure case under extreme conditions.

4.5. Impulse Response with Anisotropic Phase Functions

We consider the Henyey-Greenstein phase function (HG), parameterized by a single parameter $g \in [-1,1]$, where $g > 0$ favors forward scattering, $g < 0$ favors backward scattering, and $g = 0$ is isotropic scattering. When tabulating the canonical impulse response coefficients (Section 4.2) we compute three separate tables, one for each of $g = \{-g_{range}, 0, g_{range}\}$, where g_{range} is an artist-chosen parameter depending on the degree of desired anisotropy. All our results use $g_{range} = 0.5$. We do this only for the directly-attenuated and single-scattered tables, as we only treat isotropic scattering for multiple-scattering (see Section 5.1).

We track both an effective scattering coefficient $\sigma_s^{\text{effective}}$ **and** an effective HG anisotropy $g^{\text{effective}}$ per ray during lighting precomputation (Section 5). Given $g^{\text{effective}}$, we linearly interpolate the ZH coefficients from the three g-dependent impulse response tables and, for multiple-scattering (Section 5.1) we adjust our query σ_s to be $(1-g)\sigma_s$ according to similarity theory [ZRB14].

This is equivalent to building a discrete function space $b_i(\theta)$ to parameterize the space of phase functions, such that any phase function can be expressed as $f_p(\theta) = \mathbf{w} \cdot \mathbf{b}(\theta) = \sum_i w_i\, b_i(\theta)$, and then computing single-scattered impulse response *matrices* $\mathbf{M}$ with ZH response coefficient columns due to each basis phase function b_i. The matrix-vector product $\mathbf{M} \cdot \mathbf{w}$ would result in the appropriate ZH impulse response function for the arbitrary phase function. This does not scale to multiple scattering events: K-scattering events requires a $(K+1)$-dimensional *tensor*, that has to be collapsed to a response vector for the given per-bounce scattering parameters along the ray. With N basis phase functions, you would require $N^{(K+1)}$ coefficients for the tensor, which is impractical.

5. Using Impulse Responses to Precompute Lighting

Given pretabulated impulse response coefficients (Section 4.2) and approaches to query them for arbitrary receiver, emitter and media configurations (Sections 4.3 to 4.5), we first detail the process of the incorporating volume-to-point effects from volume attenuated and single-scattered contributions due to direct- and indirect-illumination from surfaces and emitters, below. We then discuss our treatment of multiple-scattering and the effects of occlusion on our approximation (Sections 5.1 and 5.2).

From Surface-only Baking to Volumetric Transport. Integrating volume-to-point interactions atop standard path-tracers using surface-transport light baking is straightforward with our method. Moreover, we effectively avoid the cost and complexity of tracing additional volumetric paths inside media and between the media and surfaces.

We trace standard surface-transport paths, but keep track of when rays enter and exit media. For any path vertex $\mathbf{x}_i$ with a ray entering or exiting media, we weigh the radiance $L(\overleftarrow{\mathbf{x}_i\mathbf{x}_{i+1}})$ to $\mathbf{x}$ from the direction towards the next path vertex $\overrightarrow{\mathbf{x}_i\mathbf{x}_{i+1}}$ by the appropriately queried ZH impulse response vector at $\mathbf{x}$: determined by tracking the (potentially heterogeneous) σ_s along the ray for the path segment $\mathbf{x}_i\mathbf{x}_{i+1}$ (Section 4.4) and computing the necessary σ-ratio and g (Sections 4.3 and 4.5). Weighing this ZH vector amounts to replacing I_0 in the L_e terms of Eqs. (14) to (16) with $L(\overleftarrow{\mathbf{x}_i\mathbf{x}_{i+1}})$.

Note that when $\mathbf{x}_{i+1}$ lies on an emitter (i.e., for explicit direct lighting connections in path tracing), the weighted ZH vector models the incident radiance due to directly attenuated light emission and single scattering from the light source; when $\mathbf{x}_{i+1}$ lies on another surface (i.e., for implicit indirect lighting connections), it models the incident radiance due to directly attenuated indirect surface radiance and single-scattered events that arise from a bounce of light from indirectly lit surfaces into the volume and then towards the receiver $\mathbf{x}$. For $\mathbf{x}_{i+1}$ on emitters, we use the point or directional lookup tables for these types of emitters, and the differential area table for area sources. For $\mathbf{x}_{i+1}$ on (indirectly-scattering) surfaces, we query the differential area table. When querying the differential area table, we need to also compute the angle between $\overrightarrow{\mathbf{x}_i\mathbf{x}_{i+1}}$ and $\vec{n}_{\mathbf{x}_{i+1}}$.

Given the queried ZH incident radiance vector $\mathbf{l}_{in}$ at $\mathbf{x}$, the final step in precomputed shading is to compute the outgoing radiance contribution at $\mathbf{x}$, i.e., in the lightmap. For lightprobes, we typically directly store the incident radiance $\mathbf{l}_{in}$ and convert to outgoing radiance when shading dynamic objects, using the same process we detail next for the lightmap shading case. Given a cosine-weighted BRDF f_r and ZH incident radiance $L_{in}(\mathbf{x}, \omega_i) = \mathbf{l}_{in} \cdot \mathbf{y}(\omega_i)$, we wish to compute the double-product integral

$$\int_{\mathcal{S}^2} f_r(\mathbf{x}, \omega, \omega_i)\, L_{in}(\mathbf{x}, \omega_i)\, \mathrm{d}\omega_i = \mathbf{f_r}(\omega) \cdot \mathbf{l}_{in} \text{ from Section 3.3}, \quad (17)$$

where are the SH coefficients of the view-evaluated cosine-weighted BRDF $\mathbf{f_r}(\omega)$ at $\mathbf{x}$, rotated to align in the local coordinate frame of the incident radiance's axis of alignment. In general, evaluating Eq. (17) can be costly for arbitrary BRDFs since their projection coefficients would have to either be pre-tabulated for many outgoing directions ω or computed on-the-fly. Moreover, the SH rotation incurs an additional cost; alternatively, we can rotate the ZH incident radiance more efficiently (Section 3.3) to the local frame at $\mathbf{x}$.

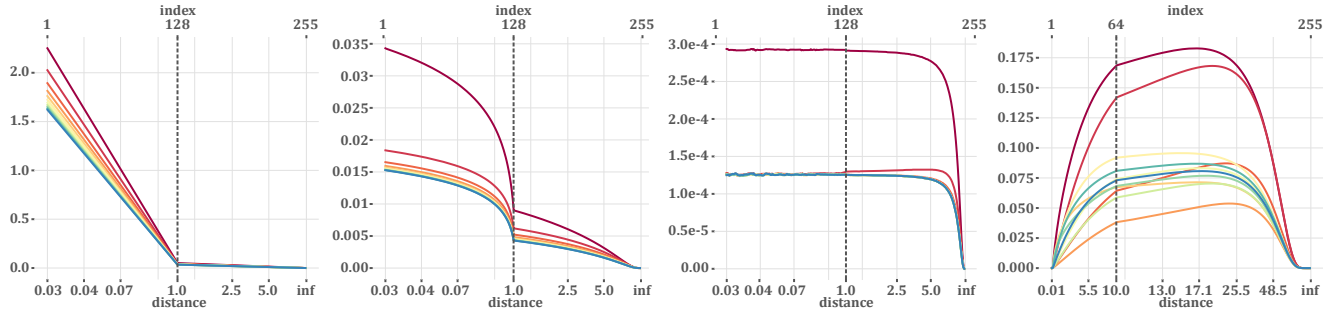

Figure 10: *Order-11 ZH response coefficients (dark red to blue) vs. canonical distances $d_{\mathbf{x},\mathbf{x}_s}$ or $d_{\mathbf{x},\mathbf{x}_{top}}$: 1-, 2- and 10-bounce response for point emitters, and 1-bounce for directional emitters. Table bin indices are on the top x-axis. Radiance tends to isotropy with higher bounces, increasing the ZH DC component (dark red) as expected. Differential area emitter plots (not shown) behave similarly to the point emitter.*

With special-case circularly symmetric BRDFs, like Lambertian or Phong, their (potentially view-evaluated) projection coefficients can be expressed in ZH along an axis of symmetry, then rotated into an appropriate frame for shading. This still requires expanding at least one of the two terms in Eq. (17) into full SH. When both terms in the double-product integrand exhibit circular symmetry, but along different axes, we develop a new fast-ZH shading formulation that avoids expanding any of the two terms into full SH (Appendix A). We use this formulation in these special-case scenarios, typically on diffuse surfaces in the lightmap where the cosine-weighted BRDF is a circularly-symmetric function about the surface $\vec{n}_{\mathbf{x}}$ normal at $\mathbf{x}$, with analytically derivable ZH coefficients [RH01].

5.1. Single- and Multiple-scattering Tables

Without absorption, we use our single-scattering tables to accelerate the computation of multiple-scattered ZH incident radiance impulse vectors: much like how indirect bounces are precomputed in surface-based PRT, where indirect ray intersections accumulate transport vector contributions from previous bounce's PRT simulation [SKS02], we can recompute the MC integral estimates of the directly-attenuated and single-scattered responses in Eqs. (14) to (16), but each time now replacing the L_e terms with the previous bounce's ZH inscattered response evaluated in the appropriate direction using the SH/ZH expansion equation in Section 3.3. Here, the same σ-ratio, heterogeneous and HG phase function remappings (Sections 4.3 to 4.5) can be applied at multiple-scattered bounces, and we can combine scattering impulse response vectors across each of these bounces into a **single lookup table**.

When $\sigma_a \neq 0$ we need to individually tabulate the multiple-scattering and, in the case of anisotropic media, we adjust isotropic scattering according to a similarity theory mapping. Specifically, we weight the ZH response of each bounce-b of multiple scattering by a factor of $\left(\sigma_s'/\sigma_t\right)^b$, as per the discussion concerning the σ-ratio adjustments in the presence of absorption in Sections 4.2 and 4.5.

We always tabulate order-11 ZH response, with 256 bins for discretized distances $d_{\mathbf{x},\mathbf{x}_{top}}$ or $d_{\mathbf{x},\mathbf{x}_s}$, and eight discrete θ angle bins. Fig. 10 illustrates the form of the 11-vector ZH response, and the tabulation ranges and steps, for our tables. As the number of scattering bounces increases the signal tends towards lower directionality (i.e., high-order coefficients lose magnitude relative to low-order ones).

5.2. Occlusion and Back-scattered Contributions

In general, light may scatter around occluders in participating media (Fig. 3) and our impulse responses do not model this effect. After analyzing the solid angle subtended by the impulse response lobes, however, we found that single-scattered response is limited to angular extents between 5-to-6°, across all configurations. This reduces the impact of ignoring occluders during pre-tabulation, and we further mitigate this by testing for occlusion with direct emission rays and only computing impulse contributions for unoccluded rays.

Another important factor to take into consideration is light that scatters volumetrically off of surfaces that are *behind* (or co-planar with) $\mathbf{x}$, before arriving back at $\mathbf{x}$. We observe empirically that this *back-scattered contribution* can be significant, and so we compensate for it during impulse response MC pre-tabulation, as follows: we slightly shift $\mathbf{x} \rightarrow \mathbf{x} + \varepsilon\vec{n}_{\mathbf{x}}$ and sample incident rays over the *entire sphere of directions* $\mathcal{S}^2$, instead of just about the hemisphere Ω at $\vec{n}_{\mathbf{x}}$, when accumulating ZH impulse response. We always set $\varepsilon = 0.125$. We artificially set the distances of these back-scattered ray intersections to 0, as if $\mathbf{x}$ had not been shifted at all. Fig. 9 illustrates the impact of the back-scattered contribution in reducing the error compared to ground truth.

6. Results & Discussion

We implement our method in a game studio's baking pipeline, supporting volume-to-surface transport for light maps, light grids and lightprobe-based meshes. Our implementation uses Intel® Embree ray tracing kernels and runs on the CPU. The runtime engine models single scattering to the eye using a modified implementation of existing work [Wro14], ignoring the effects of media between the light source and froxel grid and only treating media back to the eye.

To our knowledge, no existing method targets our problem: accelerating the computation of volume-to-surface transport contributions for the aforementioned baking targets. We compare our method against the most commonly employed solutions in the industry, in descending order of use: to ignore the effects of media scattering on surfaces, to add direct-attenuation volume contributions only (i.e., media shadowing) and to brute-force baked volume-to-surface transport. For large-scale production scenes with both homogeneous air density and dozens of heterogeneous media volumes, adding volume-to-surface transport with our method, including the cost

Figure 11: *In-game screenshot with heterogeneous media (a). Lightmaps and probe-illuminated meshes without (b) and with (c) media-scattered lighting. Ignoring volume-to-surface altogether (d) misses key transport paths. Our method efficiently captures volumetric single- and multi-scattering light from emitters and surfaces (e). ©Activision Publishing, Inc.*

of directly-attenuated lighting, incurs a modest ~4 to 15% overhead atop *surface-only baking*. We benchmark using high- and lower-quality tracing settings to differentiate between final bake and iterative workflow modes. For the same scene (shown from various viewpoints in Figs. 1, 11 and 12), baking takes **10.5 / 3.7 minutes** without and **12 / 3.8 minutes** with our volume-to-surface transport, for **high-/low-**quality: corresponding to a **14.6% / 4.7%** overhead. Runtime performance remains unchanged.

Limitations & Future Work. We discuss the approximation in single- and multiple-scattering (but not directly attenuated radiance) due to our treatment of heterogeneous (Section 4.4), and addressing this limitation can further improve our accuracy (see Fig. 8).

Ignoring the effects of light scattering around occluders is reasonable for single-scattering, due to the limited angular extent of its impulse response; the impact on multiple-scattering, however, can be larger (i.e.,around a corners or small occluders; Fig. 3, bottom right). The amount of energy in multiple-scattered events diminishes quickly in thin media [KMM*17]. Better modeling local occlusion to compensate for this lost energy is an avenue we leave to future work. We also leave the treatment of glossy materials to future work, likely necessitating alternative representations (i.e., spherical radial basis functions) that scale more compactly to higher-order angular variation.

7. Conclusion

We present a method for incorporating volumetric transport during lighting precomputation for interactive graphics applications. Our results agree closely to ground truth path tracing, at a fraction of the cost: we incorporate volumetrically-attenuated emission, single- and multiple-scattering from volumetric media with a modest overhead of 4-15% compared to *existing surface-only precomputation times*. It is straightforward to implement in existing pipelines, and has been implemented at a large gaming studio. Our impulse response tables are compact and need to be precomputed only once.

Our spherical impulse response analysis of directly-attenuated and single-scattered airlight integrals validates the utility of ZH representations for this problem. We generalize our canonical (i.e., unit homogeneous, isotropic) impulse response analysis to account for media with arbitrary scattering, extinction and absorption coefficients, arbitrary phase functions, and heterogeneity. Our algorithm is parallelizable and, as we the target static lighting pipeline, our method requires no added runtime cost. Despite this, the subtle but important lighting cues that we account for with volume-to-surface/point transport allows digital content artists to better employ media when developing complex and realistic virtual environments.

8. Acknowledgements

We thank the anonymous reviewers for their suggestions on improving our exposition, as well as Michal Drobot, Michał Iwanicki, Kyle McKisic and Ari Silvennoinen for insightful discussions. We thank Wenzel Jakob for the Mitsuba renderer [Jak13] we extended to generate the images in Figs. 8 and 9, as well as blendswap.com artist Wig42 and Benedikt Bitterli for the staircase scene in Fig. 8. This work was supported by funding from Activision and an NSERC Discovery Grant (RGPIN-2018-05669).

Appendix A: Efficient Cosine Double-product Integration

Computing Eq. (17) with a diffuse BRDF is required to convert the ZH incident radiance from our impulse response into outgoing radiance. We derive an efficient $O(N)$ integration scheme that avoids the $O(N^2)$ computation of SH coefficients for either of these (independently symmetric) factors.

We use the *SH addition theorem*, that expresses ZH rotated along any axis $\bar{\omega}$ as a product of its canonical z-axis oriented coefficients n_l^* and the value of SH basis functions evaluated at $\bar{\omega}$: $z_l^0(\omega \to \bar{\omega}) = \sum_l n_l^* \sum_m y_l^m(\omega) y_l^m(\bar{\omega})$.

The symmetry of the two ZH expansions allows a coordinate system where the first function $g(\theta)$ (with coefficients g_l) is aligned about a canonical axis ω_c and the second $h(\theta)$ (with h_l) is related to g by the angular difference in their axes. Reconstructing g along this axis is non-zero only for the $m = 0$ ZH, and h need only be evaluated at these coefficients, simplifying its reconstruction to $h(\omega) = \sum_l g_l y_l^0(\omega_c) h_l y_l^0(\omega \cdot \bar{\omega})$, where ω is ω_c in the transformed space and $(\omega \cdot \bar{\omega})$ is the cosine of the angle between the axes.

Figure 12: *Homogeneous media lit by point and directional emitters. Final renderings (a, b) include scattering towards the eye, lightmap visualizations (c, d) only include indirect light. Only treating directly-attenuated light transport (a, c) fails to capture important volume-to-surface scattering effects. Our method (b, d) accounts for these with minimal overhead, avoiding costly multi-bounce media path tracing.* ©*Activision Publishing, Inc.*

SH reconstruction is

$$h(\omega) = \sum_l \frac{g_l}{y_l^0(\omega_c)} y_l^0(\omega_c) \frac{h_l}{y_l^0(\omega_c)} y_l^0(\omega \cdot \bar{\omega}) = \sum_l g_l h_l y_l^0(\omega \cdot \bar{\omega}) / y_l^0(\omega_c).$$

References

[Abr00] ABRASH, M. "Quake's lighting model: surface caching". (2000). URL: https://www.bluesnews.com/abrash/chap68.shtml 2.

[BJ17] BITTERLI, B. and JAROSZ, W. "Beyond points and beams: higher-dimensional photon samples for volumetric light transport". *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 36.4 (July 2017). DOI: 10/gfznbr 2.

[BXH*18] BELCOUR, L., XIE, G., HERY, C., MEYER, M., JAROSZ, W., and NOWROUZEZAHRAI, D. "Integrating clipped spherical harmonics expansions". *ACM Transactions on Graphics* 37.2 (Mar. 2018). DOI: 10/gd52pf 5.

[Cha18] CHAN, D. "Material advances in Call of Duty: WWII". *Advances in Real-Time Rendering in Games*. ACM SIGGRAPH Course Notes. 2018. DOI: 10/gf3tbf 1.

[Cha60] CHANDRASEKHAR, S. *Radiative Transfer*. Dover Books on Advanced Mathematics. NY: Dover Publications, 1960 2.

[Che08] CHEN, H. "Lighting and material of Halo 3". *Game Developers Conference*. 2008 1, 2.

[CJ16] CHRISTENSEN, P. H. and JAROSZ, W. "The path to path-traced movies". *Foundations and Trends® in Computer Graphics and Vision* 10.2 (Oct. 2016). DOI: 10/gfjwjc 2.

[CNS*11] CRASSIN, C., NEYRET, F., SAINZ, M., GREEN, S., and EISEMANN, E. "Interactive indirect illumination using voxel cone tracing". *Computer Graphics Forum (Proceedings of Pacific Graphics)* 30.7 (Sept. 2011). DOI: 10/fqnc9v 2.

[dEon16] D'EON, E. *A Hitchhiker's Guide to Multiple Scattering*. Self-published, 2016. URL: http://www.eugenedeon.com/hitchhikers 2.

[DJBJ19] DENG, X., JIAO, S., BITTERLI, B., and JAROSZ, W. "Photon surfaces for robust, unbiased volumetric density estimation". *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 38.4 (July 2019). DOI: 10.1145/3306346.3323041 2.

[ERDS14] ELEK, O., RITSCHEL, T., DACHSBACHER, C., and SEIDEL, H.-P. "Principal-ordinates propagation for real-time rendering of participating media". *Computers & Graphics* 45 (Dec. 1, 2014). DOI: 10/gfzq7q 2.

[FWKH17] FONG, J., WRENNINGE, M., KULLA, C., and HABEL, R. "Production volume rendering". *ACM SIGGRAPH Course Notes*. (Los Angeles, California). New York, NY, USA: ACM Press, 2017. DOI: 10/gfzp5r 2.

[GSHG98] GREGER, G., SHIRLEY, P., HUBBARD, P. M., and GREENBERG, D. P. "The Irradiance Volume". *IEEE Computer Graphics & Applications* 18.2 (March/April 1998). DOI: 10/ckjbb8 2.

[Hil15] HILLAIRE, S. "Towards unified and physically-based volumetric lighting in Frostbite". *Advances in Real-Time Rendering in 3D Graphics and Games*. ACM SIGGRAPH Course Notes. 2015. DOI: 10/gf3s6n 1.

[Hil18] HILLAIRE, S. "Real-time raytracing for interactive global illumination workflows in Frostbite". *Game Developers Conference*. 2018 2.

[HPB06] HAŠAN, M., PELLACINI, F., and BALA, K. "Direct-to-indirect transfer for cinematic relighting". *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 25.3 (July 1, 2006). DOI: 10/cqgn89 2.

[ICG86] IMMEL, D. S., COHEN, M. F., and GREENBERG, D. P. "A radiosity method for non-diffuse environments". *Computer Graphics (Proceedings of SIGGRAPH)* 20.4 (Aug. 1986). DOI: 10/dmjm9t 2.

[IS17a] IWANICKI, M. and SLOAN, P.-P. "Ambient dice". *Proceedings of EGSR (Experimental Ideas & Implementations)*. 2017. DOI: 10/gf3s9k 2.

[IS17b] IWANICKI, M. and SLOAN, P.-P. "Precomputed lighting in Call of Duty: Infinite Warfare". *Advances in Real-Time Rendering in Games*. ACM SIGGRAPH Course Notes. 2017. DOI: 10/gf3tbc 1, 2.

[Jak13] JAKOB, W. *Mitsuba Renderer*. 2013. URL: http://www.mitsuba-renderer.org 6, 9.

[JC98] JENSEN, H. W. and CHRISTENSEN, P. H. "Efficient simulation of light transport in scenes with participating media using photon maps". *Annual Conference Series (Proceedings of SIGGRAPH)*. ACM Press, July 1998. DOI: 10/b64p36 2.

[JCJ09] JAROSZ, W., CARR, N. A., and JENSEN, H. W. "Importance sampling spherical harmonics". *Computer Graphics Forum (Proceedings of Eurographics)* 28.2 (Apr. 1, 2009). DOI: 10/b523xg 5.

[JDZJ08] JAROSZ, W., DONNER, C., ZWICKER, M., and JENSEN, H. W. "Radiance caching for participating media". *ACM Transactions on Graphics* 27.1 (Mar. 1, 2008). DOI: 10/cwnw78 2.

[JMLH01] JENSEN, H. W., MARSCHNER, S. R., LEVOY, M., and HANRAHAN, P. "A practical model for subsurface light transport". *Annual Conference Series (Proceedings of SIGGRAPH)*. ACM Press, Aug. 2001. DOI: 10/chdvp7 2.

[JNSJ11] JAROSZ, W., NOWROUZEZAHRAI, D., SADEGHI, I., and JENSEN, H. W. "A comprehensive theory of volumetric radiance estimation using photon points and beams". *ACM Transactions on Graphics* 30.1 (Jan. 1, 2011). DOI: 10/fcdh2f 2.

[JNT*11] JAROSZ, W., NOWROUZEZAHRAI, D., THOMAS, R., SLOAN, P.-P., and ZWICKER, M. "Progressive photon beams". *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 30.6 (Dec. 2011). DOI: 10/fn5xzj 2.

[JZJ08a] JAROSZ, W., ZWICKER, M., and JENSEN, H. W. "Irradiance gradients in the presence of participating media and occlusions". *Computer Graphics Forum (Proceedings of the Eurographics Symposium on Rendering)* 27.4 (June 2008). DOI: 10/bg8nww 2.

[JZJ08b] JAROSZ, W., ZWICKER, M., and JENSEN, H. W. "The beam radiance estimate for volumetric photon mapping". *Computer Graphics Forum (Proceedings of Eurographics)* 27.2 (Apr. 2008). DOI: 10/bjsfsx 2.

[Kaj86] KAJIYA, J. T. "The rendering equation". *Computer Graphics (Proceedings of SIGGRAPH)* 20.4 (Aug. 1986). DOI: 10/cvf53j 2.

[KBS11] KAVAN, L., BARGTEIL, A. W., and SLOAN, P.-P. "Least squares vertex baking". *Computer Graphics Forum (Proceedings of the Eurographics Symposium on Rendering)* 30.4 (June 1, 2011). DOI: 10/dg89k6 2.

[KD10] KAPLANYAN, A. and DACHSBACHER, C. "Cascaded light propagation volumes for real-time indirect illumination". *Proceedings of the Symposium on Interactive 3D Graphics and Games*. I3D '10. Washington, D.C.: ACM, 2010. DOI: 10.1145/1730804.1730821 2.

[KF12] KULLA, C. and FAJARDO, M. "Importance sampling techniques for path tracing in participating media". *Computer Graphics Forum (Proceedings of the Eurographics Symposium on Rendering)* 31.4 (June 2012). DOI: 10/f35f4k 2, 5.

[KMM*17] KALLWEIT, S., MÜLLER, T., MCWILLIAMS, B., GROSS, M., and NOVÁK, J. "Deep scattering: rendering atmospheric clouds with radiance-predicting neural networks". *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 36.6 (Nov. 2017). DOI: 10/gfzq7g 9.

[Leh07] LEHTINEN, J. "A framework for precomputed and captured light transport". *ACM Transactions on Graphics* 26.4 (Oct. 2007). DOI: 10/cmmfvq 2.

[LWDB10] LAURIJSSEN, J., WANG, R., DUTRÉ, P., and BROWN, B. J. "Fast estimation and rendering of indirect highlights". *Computer Graphics Forum (Proceedings of the Eurographics Symposium on Rendering)* 29.4 (June 1, 2010). DOI: 10/ctkmnv 2.

[McT04] MCTAGGART, G. "Half-Life 2 source shading". *Game Developers Conference*. 2004 1, 2.

[MJJG18] MARCO, J., JARABO, A., JAROSZ, W., and GUTIERREZ, D. "Second-order occlusion-aware volumetric radiance caching". *ACM Transactions on Graphics* 37.2 (July 2, 2018). DOI: 10/gdv86k 2.

[MPG*16] MÜLLER, T., PAPAS, M., GROSS, M., JAROSZ, W., and NOVÁK, J. "Efficient rendering of heterogeneous polydisperse granular media". *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 35.6 (Nov. 2016). DOI: 10/f9cm65 2.

[MWM07] MOON, J. T., WALTER, B., and MARSCHNER, S. R. "Rendering discrete random media using precomputed scattering solutions". *Rendering Techniques (Proceedings of the Eurographics Symposium on Rendering)*. (Grenoble, France). Eurographics Association, June 2007. DOI: 10/gfzp5n 2.

[NGHJ18] NOVÁK, J., GEORGIEV, I., HANIKA, J., and JAROSZ, W. "Monte Carlo methods for volumetric light transport simulation". *Computer Graphics Forum (Proceedings of Eurographics State of the Art Reports)* 37.2 (May 1, 2018). DOI: 10/gd2jqq 2.

[NP15] NEUBELT, D. and PETTINEO, M. "Advanced lighting R&D at Ready At Dawn Studios". *Physically Based Shading in Theory and Practice*. ACM SIGGRAPH Course Notes. 2015. DOI: 10/gf3s6p 2.

[ODo18] O'DONNELL, Y. "Precomputed global illumination in Frostbite". *Game Developers Conference*. 2018 1, 2.

[PJH16] PHARR, M., JAKOB, W., and HUMPHREYS, G. *Physically Based Rendering: From Theory to Implementation*. 3rd. Cambridge, MA: Morgan Kaufmann, 2016 2, 5.

[RH01] RAMAMOORTHI, R. and HANRAHAN, P. "An efficient representation for irradiance environment maps". *Annual Conference Series (Proceedings of SIGGRAPH)*. ACM Press, 2001. DOI: 10/bj3n3b 2, 8.

[SHHS03] SLOAN, P.-P., HALL, J., HART, J., and SNYDER, J. "Clustered principal components for precomputed radiance transfer". *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 22.3 (July 1, 2003). DOI: 10/dbvt9z 2.

[SJJ12] SCHWARZHAUPT, J., JENSEN, H. W., and JAROSZ, W. "Practical Hessian-based error control for irradiance caching". *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 31.6 (Nov. 2012). DOI: 10/gbb6n4 2.

[SKS02] SLOAN, P.-P., KAUTZ, J., and SNYDER, J. "Precomputed radiance transfer for real-time rendering in dynamic, low-frequency lighting environments". *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 21.3 (2002). DOI: 10/fgq3kn 2, 8.

[SL17] SILVENNOINEN, A. and LEHTINEN, J. "Real-time global illumination by precomputed local reconstruction from sparse radiance probes". *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 36.6 (Nov. 2017). DOI: 10/gcqbvn 2.

[Slo08] SLOAN, P.-P. "Stupid spherical harmonics (SH) tricks". *Game Developers Conference*. Feb. 2008 2.

[SLS05] SLOAN, P.-P., LUNA, B., and SNYDER, J. "Local, deformable precomputed radiance transfer". *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 24.3 (July 31, 2005). DOI: 10/bvb8qw 2, 4, 5.

[SRNN05] SUN, B., RAMAMOORTHI, R., NARASIMHAN, S. G., and NAYAR, S. K. "A practical analytic single scattering model for real time rendering". *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 24.3 (July 2005). DOI: 10/fgnbqt 2.

[WB11] WRENNINGE, M. and BIN ZAFAR, N. "Production volume rendering 1: Fundamentals". *ACM SIGGRAPH Course Notes*. ACM Press, 2011 2.

[Wro14] WRONSKI, B. "Volumetric fog: unified compute shader-based solution to atmospheric scattering". *Advances in Real-Time Rendering in 3D Graphics and Games*. ACM SIGGRAPH Course Notes. 2014. DOI: 10/gf3s6q 1, 2, 8.

[ZDM13] ZHANG, Y., DONG, Z., and MA, K.-L. "Real-time volume rendering in dynamic lighting environments using precomputed photon mapping". *IEEE Transactions on Visualization and Computer Graphics* 19.8 (Aug. 2013). DOI: 10.1109/TVCG.2013.17 2.

[ZRB14] ZHAO, S., RAMAMOORTHI, R., and BALA, K. "High-order similarity relations in radiative transfer". *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 33.4 (July 2014). DOI: 10/f6cz6p 7.