

Progressive Virtual Beam Lights

Jan Novák^{1,2} Derek Nowrouzezahrai^{1,3} Carsten Dachsbacher² Wojciech Jarosz¹
¹Disney Research Zürich ²Karlsruhe Institute of Technology ³Université de Montréal




Figure 1: Indirect illumination on surfaces and in the medium. Our method (right) significantly reduces singularity artifacts compared to the virtual ray lights method (left). We effectively eliminate singularities in four indirect light transport paths (bottom) using novel importance sampling schemes and a new lighting primitive compatible with progressive rendering.

Abstract

A recent technique that forms virtual ray lights (VRLs) from path segments in media, reduces the artifacts common to VPL approaches in participating media, however, distracting singularities still remain. We present Virtual Beam Lights (VBLs), a progressive many-lights algorithm for rendering complex indirect transport paths in, from, and to media. VBLs are efficient and can handle heterogeneous media, anisotropic scattering, and moderately glossy surfaces, while provably converging to ground truth. We inflate ray lights into beam lights with finite thicknesses to eliminate the remaining singularities. Furthermore, we devise several practical schemes for importance sampling the various transport contributions between camera rays, light rays, and surface points. VBLs produce artifact-free images faster than VRLs, especially when glossy surfaces and/or anisotropic phase functions are present. Lastly, we employ a progressive thickness reduction scheme for VBLs in order to render results that converge to ground truth.

Categories and Subject Descriptors (according to ACM CCS): I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Ray Tracing; I.6.8 [Simulation and Modeling]: Simulation—Monte Carlo

1. Introduction and Previous Work

Indirect lighting on surfaces and in scattering media greatly influences the appearance of natural phenomena like clouds, fog, and most non-metallic surfaces. In graphics, many classes of algorithms have proposed solutions or approximations to the radiative transport [Cha60] and the rendering equations [Kaj86] which govern this behavior. Cezaro et al. [CPCP*05] provide a recent overview of available techniques. We propose a novel many-light approach for indirect lighting in, from, and to volumetric media while robustly accounting for glossy scattering in the media and on surfaces.

Instant radiosity [Kel97] has led to a wide range of many-light algorithms that efficiently simulate indirect illumination (see [RDGK12] for an overview). A fundamental drawback of most of these methods are singularities, which become increasingly problematic in environments with high-frequency scattering functions (e.g. glossy BRDFs, anisotropic phase functions). We improve virtual ray lights (VRLs) [NNDJ12], which represent transport in media using the segments (and not vertices) of random-walk paths. VRLs sample light in the media more densely than VPLs (see Figure 2) and provably reduce – but do not entirely eliminate – singularities.




Figure 2: Left: we are inspired by VSLs (bottom), which expand VPLs (top) into spheres with finite solid angle. Right: we similarly expand VRLs (top) into VBLs (bottom) with finite thickness. This removes singularities and reduces render time for indirect lighting to and from surfaces and media.

Singularities are normally handled with clamping, but this causes bias due to energy loss. Bias compensation [KK06, RSK08, END10] can restore the missing energy but can be expensive, especially in glossy environments. Hašan et al. [HKWB09] proposed an alternative to clamping for surface illumination; by blurring the contribution of VPLs into virtual spherical lights (VSLs, see Figure 2), infinite singularities no longer appear. Effectively, VSLs act as an explicit final gather over standard photon mapping [Jen01].

Contributions. We address the shortcomings of the VRL method: eliminating singularities and explicitly addressing the problematic transport between (potentially glossy) surfaces and media. At a high-level, we extend the benefits of VSLs to participating media and VRLs. Our core idea is conceptually simple but practical: by carefully blurring VRLs into *virtual beam lights* (VBLs, see Figure 2) we eliminate singularities, even in the presence of glossy scattering in and between media and surfaces.

Though VBLs introduce some bias, a variety of related techniques [HJ09, KZ11, DKL10] have recently proposed progressive formulations to produce convergent results while maintaining a fixed memory footprint. As we will show, our approach can also be viewed as an explicit final gather over progressive photon beams [JNT*11], providing a simple progressive algorithm that allows for fast preview while converging to a high-quality result in the limit.

Overview. We briefly review radiative transport for surfaces and volumes and, as our approach extends VRLs, we also summarize the VRL method (derivations, proofs of correctness, and more details can be found in the VRL paper [NNDJ12]). Afterwards, we explain the details of our VBL approach, showing how it provides a more complete solution while avoiding issues with the VRL formulation.

2. Radiative Transport and Virtual Ray Lights

Radiance travelling in direction $\vec{\omega}$ towards a point \mathbf{x} can be expressed as a sum of surface L_s and media radiance L_m ,

$$L(\mathbf{x}, \vec{\omega}) = T_r(s) L_s(\mathbf{x}_s, \vec{\omega}) + L_m(\mathbf{x}, \vec{\omega}), \quad (1)$$

where $\mathbf{x}_s = \mathbf{x} - s\vec{\omega}$ is the nearest surface point to \mathbf{x} from direction $\vec{\omega}$. For brevity, we use a simplified, single-parameter notation for transmittance $T_r(s) = e^{-\int_0^s \sigma_t(\mathbf{x}-t\vec{\omega}) dt}$ where σ_t is the (spatially varying) extinction coefficient. Surface radiance L_s is defined by the rendering equation [Kaj86]:

$$L_s(\mathbf{x}_s, \vec{\omega}) = L_e + \int_{\Omega} f_r(\vec{\omega}, \vec{\omega}') L(\mathbf{x}_s, \vec{\omega}') d\vec{\omega}', \quad (2)$$

where L_e is the emitted radiance on light sources, and f_r is the cosine-weighted BRDF.

Media radiance L_m integrates the product of the scattering coefficient σ_s , transmittance, and lighting between \mathbf{x} and \mathbf{x}_s :

$$L_m(\mathbf{x}, \vec{\omega}) = \int_0^s \sigma_s(u) T_r(u) \int_{\Omega_{4\pi}} f_s(\vec{\omega}, \vec{\omega}') L(\mathbf{x}_u, \vec{\omega}') d\vec{\omega}' du, \quad (3)$$

where $\mathbf{x}_u = \mathbf{x} - u\vec{\omega}$ parametrizes the integration along the ray. The inner integral computes the in-scattered radiance by recursively weighting the radiance incident from all directions $\Omega_{4\pi}$ by the phase function f_s .

The radiance in Equations (2) and (3) can be further decomposed based on whether the incident light arrives from surfaces, from media, or directly from light sources:

$$L_s = L_s^s + L_s^m + L_s^l \quad \text{and} \quad L_m = L_m^s + L_m^m + L_m^l, \quad (4)$$

where subscripts denote the location the radiance is estimated, and superscripts denote the source of the radiance. L_s^l and L_m^l are the radiance arriving directly from light sources and are not considered further. The others are: surface-to-surface (L_s^s), media-to-surface (L_s^m), surface-to-media (L_m^s) and media-to-media (L_m^m) transport paths.

Virtual Ray Lights. The VRL method approximates Equation (1) by first performing a random-walk photon simulation and converting the segments of the random-walk paths into linear light sources. During rendering, this collection of virtual ray lights is used to estimate the L_s^m and L_m^m indirect lighting components. The two remaining transport types (L_s^s and L_m^s) are not represented by VRLs (but by their endpoints) and are handled with other algorithms (e.g., VPLs).

For each VRL, L_s^m in Equation (4) is approximated as:

$$L_s^m \approx \Phi \int_0^t \frac{\sigma_s(v) f_r(s) f_s(v) T_r(v) T_r(s, v) V(s, v)}{w(s, v)^2} dv, \quad (5)$$

where v parametrizes positions along the VRL of length t and power Φ . $V(s, v)$ and $w(s, v)$ are the binary visibility and distance between the point s at the end of a camera ray and a point v along the VRL (we omit a few parameters for brevity).

The L_m^m component of Equation (3) is estimated as a double

integral along the camera ray of length s and VRL:

$$L_m^m \approx \Phi \int_0^s \int_0^t \frac{\sigma_s(u) \sigma_s(v) f_s(u) f_s(v) T_r(u) T_r(v) T_r(u, v) V}{w(u, v)^2} dv du. \quad (6)$$

Denoting the integrands in Equations (5) and (6) as $g_s^m(v)$ and $g_m^m(u, v)$, these integrals can be estimated with unbiased Monte Carlo integration as:

$$L_s^m \approx \frac{1}{N} \sum_{i=1}^N \frac{g_s^m(v_i)}{\text{pdf}(v_i)} \quad \text{and} \quad L_m^m \approx \frac{1}{N} \sum_{i=1}^N \frac{g_m^m(u_i, v_i)}{\text{pdf}(u_i, v_i)}, \quad (7)$$

where $\text{pdf}(v_i)$ is the probability of choosing a location along the VRL for a given surface point, and $\text{pdf}(u_i, v_i)$ is the probability of choosing points on the camera ray and VRL (u_i, v_i) . A major contribution of Novák et al. [NNDJ12] is a method of importance sampling these estimators according to the phase functions and inverse-squared distance terms. We directly improve these two transport types and explicitly address the remaining transport that VRLs do not consider (L_m^s and L_s^s), all within a progressive rendering framework.

3. Virtual Beam Lights

Though VRLs reduce VPL singularities by distributing energy along the length of light rays, many distracting singularities remain (especially for transport between surfaces and media where only a single, and not double, integration is performed). Clamping, however, results in energy loss and modified material appearance (Figure 3). Even with perfect importance sampling of Equation (7), a large number of VRLs may be required to produce artifact-free results.




Figure 3: Equal-Virtual-Light comparison between unclamped/clamped VRLs/VPLs [NNDJ12] (left), VBLs/VSLs (right), and reference (far-right) on different transports paths.




Figure 4: Visualization of our approach for each transport type. VBLs are illustrated as swept VSLs for L_s^m and L_m^m .

To overcome this problem, we begin by blurring VRLs into volumetric VBLs. Conceptually, instead of integrating point-to-point transfer with VRLs, we inflate all points along a VRL into volumetric spheres (thus inflating the entire ray to a volumetric beam). Moreover, for surface-to-media transport, we formulate an analogous solution by integrating the contribution of a surface VSL along camera rays through the medium. To make this approach efficient, the modified transport equations require new importance sampling strategies.

3.1. Media-to-Surface Transport

We obtain the L_s^m contribution to a surface point \mathbf{x}_s due to a VBL by expanding every point along a VRL into a VSL (see Figure 4). This modifies Equation (5) to:

$$L_s^m(\mathbf{x}_s, \vec{\omega}) \approx \int_0^t \sigma_s(v) T_r(v) T_r(s, v) L_s^{\text{vsl}}(s, v) dv, \quad (8)$$

where L_s^{vsl} integrates the scattering functions over the solid angle Ω_{vsl} of the VSL centered at v from the surface at s :

$$L_s^{\text{vsl}}(s, v) = \frac{\Phi}{\pi R^2} V(s, v) \int_{\Omega_{\text{vsl}}} f_r(\vec{\omega}, \vec{\omega}') f_s(\vec{\omega}_v, \vec{\omega}') d\vec{\omega}', \quad (9)$$

where $\vec{\omega}_v$ is the VRL's direction. Equation (9) can be solved using the beam radiance estimate [JZJ08], which interprets the volumetric sphere as a disc, of radius R , that always faces the gather direction $\vec{\omega}'$ (see Figure 5, left). Figure 6 plots the contributions of a VRL (orange) and a VBL (blue) to a shading point, and also illustrates how sampling in the angular (instead of parametric) domain reduces high-frequency variations while implicitly incorporating the inverse squared distance fall-off.

In order to importance sample the location v on the VRL (i.e. to sample the VSL location) we construct a piecewise-linear PDF in the angular domain about \mathbf{x}_s , much like Novák et al. [NNDJ12]. Figure 6 visualizes the piecewise-linear PDF proposed by Novák et al. This matches the VRL response quite well, but unfortunately does not match the VBL response well, which motivates our new importance sampling technique.




Figure 5: Evaluating the volumetric VSL using a BRE final gather and estimating the max BRDF angle $\tilde{\omega}'_{max}$.

At each vertex of this PDF we would ideally evaluate the contribution of a VSL placed at that location. However, integrating the contribution (e.g. with Monte Carlo integration of Equation (9) over many directions) would be very costly and noisy due to possibly glossy BRDFs and phase functions.

Instead, we observe that while *overestimating* the integrand response in the PDF is suboptimal, it does not lead to nearly as much variance as *underestimation*. At each PDF vertex, we therefore quickly approximate the maximum contribution of the integrand's terms within Ω_{vsl} . Thus, for every PDF vertex, we (see Figure 5, right):

- find the direction $\tilde{\omega}_{max}$ where incident light at the surface contributes most to the reflection according to the BRDF,
- find the direction $\tilde{\omega}'_{max}$ with least rotation to $\tilde{\omega}_{max}$ on the cone of directions subtended by the VSL, and
- evaluate the product of BRDF and phase function at $\tilde{\omega}'_{max}$.

Note that our PDF implicitly accounts for inverse squared fall-off in Equation (8) by sampling in the angular domain, and explicitly for the BRDF's max within Ω_{vsl} (Equation (9)). Finding $\tilde{\omega}_{max}$ is relatively simple for symmetric BRDFs (e.g. Phong, Ward); we discuss arbitrary BRDFs in Section 6.

Computing light transport during rendering now becomes straightforward: we sample a point along the VRL according to our constructed PDF, inflate the point into a VSL, and numerically integrate Equation (9) using the beam radiance estimate [JZJ08] as illustrated in Figure 5 (left).

3.2. Media-to-Media Transport

For L_m^m (Figure 4, bottom left) we similarly inflate points along the VRL into VSLs, resulting in the following modification to Equation (6):

$$L_m^m \approx \int_0^s \int_0^t \sigma_s(u) \sigma_s(v) T_r(u) T_r(v) T_r(u, v) L_m^{vsl}(u, v) dv du, \quad (10)$$

where $L_m^{vsl}(u, v)$ is defined analogously to $L_s^{vsl}(s, v)$ but using two phase functions instead of a phase function and a BRDF.

For this transport type we extend our procedure in Section 3.1 by wrapping it into a two-stage method to account for the 2D integration domain. We first assume the camera ray is an infinite line and importance sample a point v_i along the




Figure 6: A VRL's (orange) and VBL's (blue) contribution to a shading point. Left: working in the angular domain smooths out the peak and implicitly accounts for inverse-squared falloff. Middle/Right: two examples of the piecewise-linear PDFs for VRLs [NNDJ12] and VBLs (dashed) compared to the actual integrand (solid).

VRL according to the inverse squared distance [NNDJ12, Section 4.1] where we place a VSL. Once we have a VSL, we proceed analogously to the media-to-surface case: we construct a piecewise-linear PDF in the angular domain about the VSL center v_i . As both locations (VSL and the point on the camera ray) are in the media, the PDF construction evaluates two phase functions instead of a phase function and a BRDF. We find $\tilde{\omega}_{max}$ based on the phase function at v_i ; for analytic models, e.g. Henyey-Greenstein, this is again trivial.

3.3. Surface-to-Media Transport

L_m^s transport (Figure 4, bottom right) is the dual of L_s^m : to sample a point on the ray we construct the PDF in the angular domain about the surface point. The surface VPL (or endpoint of the VRL) is inflated to a VSL and its contribution to the camera ray location is evaluated numerically.

3.4. Surface-to-Surface Transport

L_s^s transport is computed with a modified application of the traditional VSL approach [HKWB09]: we compute the transport progressively (see Section 4) with a shrinking VSL radius, resulting in converged results over many passes.

4. Algorithm

We implemented our algorithm in a hybrid CPU-GPU rendering system where CPU ray-tracing is used to trace path-segments from light sources through the scene (scattering in the media and at surfaces) and reflected/refracted camera rays. We use a GPU ray-tracing kernel [AL09] to shoot visibility rays for all of our estimates.

Path segments traced from light sources form our VBLs and their endpoints on surfaces yield VSLs and surface photons. The latter are used for surface caustics using progressive photon mapping (PPM) [HOJ08, KZ11]; all other indirect illumination on surfaces (from surfaces) is computed with VSLs [HKWB09] using our progressive estimation. We evaluate indirect surface lighting from scattered media radiance with Equation (8) at each shading point and from each VBL.

We compute multiple-scattering for each camera ray and VBL with Equation (10) and light scattered from surfaces back to the media with the dual to Equation (8) (as described in Section 3.3). We handle single-scattering and volume caustics using progressive photon beams (PPB) [JNT*11].

4.1. Progressive Rendering

Our complete rendering algorithm is contained within a progressive estimation framework. We render multiple independent passes, the running average of which is displayed as the rendered image.

Inflating VPLs into VSLs and VRLs into VBLs introduces bias by blurring out the illumination and scattering functions. Fortunately, each of these transport types can be viewed as an explicit final gather over either progressive photon beams [JNT*11], or progressive photon mapping [HOJ08, KZ11]. We can therefore rely on progressive radius reduction to ensure that both bias and variance converge to zero in the limit. In our context, the radius reduction for all four types of light transport corresponds to a 2D blur (surface VSLs are blurred into discs, and volumetric VSLs are integrated using the 2D blurring of the BRE). We use the improved reduction formula from Jarosz et al. [JNT*11] (which minimizes the impact of the number of photons per pass M on the final result) applied to the *squared* radius due to the 2D blurring of VSLs and VBLs. The squared radius in the i^{th} pass is:

$$R_i^2 = R_0^2 \left(\prod_{k=1}^{M_i-1} \frac{k + \alpha}{k} \right) \frac{1}{M_i}, \quad (11)$$

where $\alpha \in (0, 1)$ is a user parameter that specifies the aggressiveness of the radius reduction. We currently set the initial radius R_0 of all VBLs and VSLs to a user-specified constant, though this could be modified to use an adaptive per-VBL/VSL radius for improved results.

Figure 7 shows the BUDDHA scene, for which we detail the progressive radius reduction using zoomed-insets in Figure 8. The insets show a small region (L_m^s on the top, L_s^m on the bottom) of sharp illumination. Though the initial estimate is somewhat blurred using VBLs and VSLs, the progressive reduction converges to the reference solution (computed using many unclamped VRLs and VPLs).

5. Results

We compare the quality and performance of our technique using progressive VBLs+VSLs to that of Novák et al. [NNDJ12] which uses unclamped VRLs+VPLs. We also compare to a clamped version of VRLs+VPLs. The full global illumination solutions for our three scenes are shown in Figure 10, where we additionally include volume and surface caustics computed using PPB and PPM. All timings were measured on an Intel Core i7 with 12 CPUs @ 3.2GHz, 24GB RAM, and an NVIDIA Quadro 6000. The CPU ray-tracer and PPM/PPB is parallelized over all CPU cores. We use a




Figure 7: An equal-time (unconverged) comparison in the BUDDHA scene between unclamped/unbiased VRLs+VSLs (left), clamped VRLs+VSLs (center), and our method (right).




Figure 8: Our progressive VBLs+VSLs converge to the reference, here computed using a large number of VRLs+VPLs.

Henyey-Greenstein (HG) phase function for all our results with anisotropic media. We use up to 100 rays to numerically integrate the solid angle of VSLs, as suggested by Hašan et al. [HKWB09]. For all our results we use $\alpha \approx 0.7$.

Figure 7 shows the BUDDHA scene (720×720 resolution) which is filled with dense media with a strongly anisotropic ($g = 0.7$) HG phase function; the statue is highly glossy with a Phong exponent of 80. The majority of the illumination in the scene is due to indirect contributions. Unclamped VRLs+VPLs result in visible singularities, and clamped



Figure 9: We isolate and visualize the four different transport paths in the CARS (left half) and SMOKYROOM (right half) scenes. We compare our VBL+VSL method (second column in each set) to unclamped VRLs+VPLs [NNDJ12] at equal-time.

VRLs+VPLs lose energy and modify material appearance (note the regions around the Buddha and the orange wall).

Figures 1 and 9 (top) show the CARS scene which consists of 2.2 million triangles (1.2 million vertices). Both cars are made of highly glossy materials (Phong BRDF with exponents for the car bodies of 100 and 60). The surrounding media is homogeneous with anisotropic scattering (HG coefficient $g = 0.25$). Note that light sources are encased in glass, i.e. all direct illumination in the scene is due to caustics. Figure 1 shows that virtually no artifacts remain visible with our method after 46.2 minutes render time (1280×720 resolution). Figure 9 shows an equal-time comparison of the individual transport components that were computed using unclamped VRLs+VPLs and our VBL+VSL method; each set of transport types is tone mapped separately to maximize visibility. This figure also shows the same components for the SMOKYROOM scene, which is filled with isotropic, heterogeneous media (computed using Perlin noise) whose density decreases with height.

6. Discussion and Future Work

In general, our method eliminates distracting artifacts faster than VRLs+VPLs, especially when glossy surfaces and/or anisotropic phase functions are present (see also Figure 3). The singularities are traded for some blurring, which is less objectionable, especially for early passes, and results in less

visual changes from reference than clamping. In the limit however, this blurring disappears completely due to progressive radius reduction (which is not currently possible with clamping). We similarly improve standard VSLs, enabling convergent results with fixed memory usage. Though the mathematical convergence rate is slower than unbiased methods [KZ11], VBLs deliver an acceptable image faster, which the user often cares about more than numerical convergence.

Arbitrary BRDFs. For BRDFs that do not exhibit perfect circular symmetry, our approach should still perform well as long as the BRDF has a single major peak from where its value falls-off. In essence, the maximum angle defines a cone of directions that is guaranteed to contain the maximum value. In the case of symmetric BRDFs, all values along this cone are equal to the maximum. For general BRDFs, a more sophisticated technique could try to solve for the unique maximum within this cone. We leave this as future work.

Perfect PDF. Our approximate PDF for the product of the two scattering functions may become slightly suboptimal in certain cases. We avoid expensive numeric construction by identifying directions where the BRDF (or the phase function at the VRL for L_m^m) reaches its maximum value. We then evaluate the product of both scattering functions assuming that all the transport occurs along that direction, using it as the value when constructing the PDF. In some cases this might not give the actual maximum of the product.



Figure 10: Full global illumination for the BUDDHA, CARS, and SMOKYROOM scenes, including the sum of the four components from Figures 7 and 9 as well as caustics computed using progressive photon beams and progressive photon mapping.

Investigating alternative approaches to efficiently construct more accurate PDFs for our estimators is an interesting area of future work. One option for circularly symmetric and monotonically decreasing functions (measured about deviations from the axis of symmetry, e.g. the HG phase function or Phong BRDF) is to also identify the maximum direction of contribution of the second scattering function. The product's maximum is guaranteed to lie on the great arc between these two directions and may be, in some cases, found analytically.

7. Conclusion

We presented VBLs, a novel many-lights algorithm and a new lighting primitive that eliminates the singularities present in existing many-light techniques for volumetric media while avoiding the downsides of clamping. The benefits of VBLs over VRLs are akin to the benefits of VSLs over VPLs, although we handle a larger range of transport scenarios. We also devised novel importance sampling schemes to explicitly sample these complex transport paths in, from, and to surfaces and media. Our method converges faster than the state-of-the-art, without distracting singularities or energy loss, in scenes with complex lighting, glossy BRDFs, anisotropic phase functions, and heterogeneous media.

We extended progressive radiance estimation to handle VBLs and counteract the bias they introduce, ensuring that we converge to the correct result with a fixed memory footprint. As a result, we generate both fast previews and converged renderings with the same algorithm.

Acknowledgements. We would like to thank Benjamin Müller for investigations in the early stages of this project.

References

- [AL09] AILA T., LAINE S.: Understanding the efficiency of ray traversal on GPUs. In *Proc. High Performance Graphics* (2009), pp. 145–149. 4
- [Cha60] CHANDRASEKAR S.: *Radiative Transfer*. Dover Publications, 1960. 1
- [CPCP*05] CEREZO E., PEREZ-CAZORLA F., PUEYO X., SERON F., SILLION F.: A survey on participating media rendering techniques. *The Visual Computer* 21 (2005), 303–328. 1
- [DKL10] DAMMERTZ H., KELLER A., LENSCH H. P. A.: Progressive point-light-based global illumination. *Computer Graphics Forum* 29, 8 (2010), 2504–2515. 2
- [END10] ENGELHARDT T., NOVÁK J., DACHSBACHER C.: *Instant Multiple Scattering for Interactive Rendering of Heterogeneous Participating Media*. Tech. rep., Karlsruhe Institute of Technology, Dec. 2010. 2
- [HJ09] HACHISUKA T., JENSEN H. W.: Stochastic progressive photon mapping. *ACM Transactions on Graphics (Proc. SIGGRAPH Asia)* 28, 5 (2009), 141:1–141:8. 2
- [HKWB09] HAŠAN M., KRÍVÁNEK J., WALTER B., BALÁ K.: Virtual spherical lights for many-light rendering of glossy scenes. *ACM Transactions on Graphics (Proc. SIGGRAPH Asia)* 28, 5 (2009), 143:1–143:6. 2, 4, 5
- [HOJ08] HACHISUKA T., OGAKI S., JENSEN H. W.: Progressive photon mapping. *ACM Transactions on Graphics (Proc. SIGGRAPH Asia)* 27, 5 (2008), 130:1–130:8. 4, 5
- [Jen01] JENSEN H. W.: *Realistic Image Synthesis Using Photon Mapping*. A. K. Peters, Ltd., Natick, MA, USA, 2001. 2
- [JNT*11] JAROSZ W., NOWROUZEZAHRAI D., THOMAS R., SLOAN P.-P., ZWICKER M.: Progressive photon beams. *ACM Transactions on Graphics (Proc. SIGGRAPH Asia)* 30, 6 (2011), 181:1–181:12. 2, 5
- [JZJ08] JAROSZ W., ZWICKER M., JENSEN H. W.: The beam radiance estimate for volumetric photon mapping. *Computer Graphics Forum (Proc. of Eurographics)* 27, 2 (2008), 557–566. 3, 4
- [Kaj86] KAJIYA J. T.: The rendering equation. In *Computer Graphics (Proc. SIGGRAPH '86)* (1986), pp. 143–150. 1, 2
- [Kel97] KELLER A.: Instant radiosity. In *SIGGRAPH* (1997), pp. 49–56. 1
- [KK06] KOLLIG T., KELLER A.: Illumination in the presence of weak singularities. In *Monte Carlo and Quasi-Monte Carlo Methods 2004*. 2006, pp. 245–257. 2
- [KZ11] KNAUS C., ZWICKER M.: Progressive photon mapping: A probabilistic approach. *ACM Transactions on Graphics* 30, 3 (2011), 25:1–25:13. 2, 4, 5, 6
- [NNDJ12] NOVÁK J., NOWROUZEZAHRAI D., DACHSBACHER C., JAROSZ W.: Virtual ray lights for rendering scenes with participating media. *ACM Transactions on Graphics (Proc. SIGGRAPH)* 29, 4 (2012). to appear. 1, 2, 3, 4, 5, 6
- [RDGK12] RITSCHEL T., DACHSBACHER C., GROSCH T., KAUTZ J.: The state of the art in interactive global illumination. *Computer Graphics Forum* 31, 1 (2012), 160–188. 1
- [RSK08] RAAB M., SEIBERT D., KELLER A.: Unbiased global illumination with participating media. In *Monte Carlo and Quasi-Monte Carlo Methods 2006*. Springer, 2008, pp. 591–606. 2