

Progressive Transient Photon Beams

Julio Marco¹ Ibón Guillén¹ Wojciech Jarosz² Diego Gutierrez¹ Adrian Jarabo¹

¹ Universidad de Zaragoza, I3A ² Dartmouth College




Figure 1: The SOCCER scene (steady-state render on the right) features complex volumetric caustics due to multiple reflections and refractions off smooth dielectrics inside the medium. We are able to efficiently render the transient light transport (left sequence) by formulating a progressive, transient form of photon beam density estimation which provably eliminates error while working within a finite memory budget. Please refer to the supplemental video for the full sequence.

Abstract

In this work we introduce a novel algorithm for transient rendering in participating media. Our method is consistent, robust, and is able to generate animations of time-resolved light transport featuring complex caustic light paths in media. We base our method on the observation that the spatial continuity provides an increased coverage of the temporal domain, and generalize photon beams to transient-state. We extend steady-state photon beam radiance estimates to include the temporal domain. Then, we develop a progressive variant of our approach which provably converges to the correct solution using finite memory by averaging independent realizations of the estimates with progressively reduced kernel bandwidths. We derive the optimal convergence rates accounting for space and time kernels, and demonstrate our method against previous consistent transient rendering methods for participating media.

CCS Concepts

•Computer Graphics → Three-dimensional graphics and realism; Raytracing; Transient rendering;

1. Introduction

The emergence of transient imaging has led to a vast number of applications in graphics and vision [JMMG17], where the ability of sensing the world at extreme high temporal resolution allows new applications such as imaging light in motion [VWJ*13], appearance capture [NZV*11], geometry reconstruction [BH04, MHM*17], or vision through media [Bus05, WJS*18] and around the corner [VWG*12, AGJ17]. Sensing through media is one of the key applications: The ability of demultiplexing light interactions in the temporal domain is a very promising approach for important practical domains such as non-invasive medical imaging, underwa-

ter vision, or autonomous driving through fog. Accurately simulating light transport could help enormously in these applications, potentially serving as a benchmark, a forward model in optimization, or as a training set for machine learning.

Transient rendering in media is, however, still challenging: The increased dimensionality (time) increases variance dramatically in Monte Carlo algorithms, potentially leading to impractical rendering times. This variance is especially harmful in media, where the signal tends to be smooth due to the low-pass filtering behavior of scattering, in both the spatial and temporal domains. One of the major drawbacks of transient rendering is that it requires much

higher sampling rates to fill up the extended temporal domain, especially when using $0D$ (photon) point samples, which are sparsely distributed across both time and space. We make the observation that $1D$ photon trajectories populate both space and time much more densely; hence, a technique based on photon beams [JNSJ11] should significantly reduce the rendering time when computing a noise-free time-resolved render, and, given its density estimation nature, it could naturally combine with the temporal domain density estimation proposed by Jarabo et al. [JMM*14].

We present a new method for transient-state rendering of participating media, that leverages the good properties of density estimation for reconstructing smooth signals. Our work improves Jarabo et al. [JMM*14] by extending *progressive photon beams* (PPB) [JNT*11] to the transient domain, and combining it with temporal density estimation for improved reconstruction in both the spatial and temporal domains. Our technique is biased but consistent, converging to the ground truth using finite memory by taking advantage on the progressive [HOJ08, KZ11] nature of density estimation. We analyze the asymptotic convergence of our proposed space-time density estimation, computing the optimal kernel reduction ratios for both domains. Finally, we demonstrate our method on a variety of scenes with complex volumetric light transport, featuring high-frequency occlusions, caustics, or glossy reflections, and show its improved performance over naively extending PPB to the transient domain.

This paper is an extension of our previous work on rendering transient volumetric light transport [MJGJ17], where we proposed a naive extension of photon beams to transient state. Here we increase the applicability of the method, by proposing a progressive version of the space-time density estimation, and rigorously analyze its convergence.

2. Related Work

Rendering participating media is a long-standing problem in computer graphics, with a vast literature on the topic. Here we focus on works related directly with the scope of the paper. For a wider overview on the field, we refer to the recent survey by Novák et al. [NGHJ18].

Photon-based Light Transport. Photon mapping [Jen01] is one of the most versatile and robust methods for rendering complex global illumination, with several extensions for making it compatible with motion blur [CJ02], adapting the distribution of photons [SJ09, GRv*16], carefully selecting the radiance estimation kernel [SJ09, KD13, JRJ11], combining it with unbiased techniques [GKDS12, HPJ12], or making it progressive for ensuring consistency within a limited memory budget [HOJ08, KZ11]. Hachisuka et al.'s [HJG*13] recent SIGGRAPH course provides an in-depth overview.

Jensen and Christensen [JC98] were the first to extend photon mapping to media, and Jarosz and colleagues [JZJ08] significantly improved its efficiency with the beam radiance estimate, which replaces repeated point queries with one “beam” query finding all photons along the entire camera ray. Jarosz et al. [JNSJ11] later applied this idea to the photon tracing process by storing full photon trajectories (photon beams), leading to a dramatic increase in

photon density for the same photon tracing step. Their progressive and hybrid counterparts [JNT*11, KGH*14] leveraged the benefits of photon beams while providing consistent solutions using finite memory. Recently, Bitterli and Jarosz [BJ17] generalized $0D$ photon points and $1D$ photon beams to even higher dimensions, proposing the use of photon planes ($2D$), volumes ($3D$) and, in theory, higher-dimensional geometries, leading to unbiased density estimation. All these works are, however, restricted to steady-state renders; we instead focus on simulating light transport in transient state.

Transient rendering. Though the transport equations [Cha60, Gla95] are time-resolved, most rendering algorithms focus on steady-state light transport. Still, several works have been proposed to deal with light transport in a time-resolved manner. In particular, most previous work on transient rendering has focused on simulating surfaces transport: Klein et al. [KPM*16] extended Smiths’ transient radiosity [SSD08] for second bounce diffuse illumination, while other work has used more general methods based on transient extensions of Monte Carlo (bidirectional) path tracing [Jar12, JMM*14, PBSC14, JA18] and photon mapping [MNJK13, OHX*14]. Several works have also dealt with time-resolved transport on the field of neutron transport [CPH53, BG70, Wil71, DM79]. Closer to our work, Ament and colleagues [ABW14] rendered transient light transport in refractive media using volumetric photon mapping, but they do not provide an efficient approach that guarantees consistency. Jarabo et al. [JMM*14] proposed a transient extension of the path integral, and introduced an efficient technique for reconstructing the temporal signal based on density estimation. They also proposed a set of techniques for sampling media interactions uniformly in time. Their method is however limited to bidirectional path tracing and photon mapping, often failing to densely populate media in the temporal domain. Finally, Bitterli [Bit16b] and Marco et al. [Mar13, MJGJ17] proposed a transient extension of the photon beams algorithm, but these approaches are not progressive, therefore not converging to the correct solution in the limit. Our work extends the latter, proposing a progressive, consistent, and robust method for rendering transient light transport. We leverage beams continuity and spatio-temporal density estimation to mitigate variance in the temporal domain, and derive the parameters for optimal convergence of the method.

3. Transient Radiative Transfer

The *radiative transfer equation* (RTE) [Cha60] models the behavior of light traveling through a medium. While the original formulation is time-resolved, its integral form used in traditional rendering ignores this temporal dependence, and computes the radiance L reaching any point \mathbf{x} from direction $\vec{\omega}$ as

$$L(\mathbf{x}, \vec{\omega}) = T_r(\mathbf{x}, \mathbf{x}_s) L_s(\mathbf{x}_s, \vec{\omega}) + \int_0^s \mu_s(\mathbf{x}_q) T_r(\mathbf{x}, \mathbf{x}_q) L_o(\mathbf{x}_q, \vec{\omega}) dq, \quad (1)$$

where $\mathbf{x}_d = \mathbf{x} - d \cdot \vec{\omega}$ is a point at distance d , μ_s is the scattering coefficient, and $T_r(\mathbf{x}, \mathbf{x}_d) = \exp(-\int_0^d \mu_t(\mathbf{x}_{d'}) dd')$ is the *transmittance* describing the fraction of photons that make it between \mathbf{x} and \mathbf{x}_d without undergoing extinction at any point $\mathbf{x}_{d'}$. The outgoing radiance L_o in direction $\vec{\omega}$ from a medium point \mathbf{x}_q at distance q is defined by the

scattering integral:

$$L_o(\mathbf{x}_q, \vec{\omega}) = L_e(\mathbf{x}_q, \vec{\omega}) + \int_{\mathcal{S}} f_s(\mathbf{x}_q, \vec{\omega}_i, \vec{\omega}) L(\mathbf{x}_q, \vec{\omega}_i) d\vec{\omega}_i, \quad (2)$$

where \mathcal{S} is the spherical domain, and f_s is the phase function. L_s is defined analogously via the *rendering equation* [Kaj86], but integrated over the hemispherical domain, and using the cosine-weighted BSDF in place of the phase function.

Transient RTE Equations 1 and 2 assume that the speed of light is infinite. However, if we want to solve the RTE at time scales comparable to the speed of light we need to incorporate the different delays affecting light. In the following we review the main practical considerations for accounting time into the integral form of the RTE for its application in transient rendering. Light takes a certain amount of time to propagate through space, and therefore light transport from a point \mathbf{x}_0 towards a point \mathbf{x}_1 does not occur immediately. In the absence of scattering effects, transport between two points \mathbf{x}_0 and \mathbf{x}_1 occurs as

$$L(\mathbf{x}_1, \vec{\omega}, t) = L(\mathbf{x}_0, -\vec{\omega}, t - \Delta t), \quad (3)$$

where Δt is the time it takes the light to go from \mathbf{x}_0 to \mathbf{x}_1 . In turn, Δt is defined by

$$\Delta t(\mathbf{x}_0 \leftrightarrow \mathbf{x}_1) = \int_{\mathbf{x}_0}^{\mathbf{x}_1} \frac{\eta(\mathbf{x})}{c} d\mathbf{x}, \quad (4)$$

where $\eta(\mathbf{x})$ is the index of refraction at a medium point \mathbf{x} and c is the speed of light in vacuum. Note that in this case light does not travel in a straight line, but by following the Eikonal equation [ABW14, GMAS05]. In a medium with a constant index of refraction $\eta(\mathbf{x}) = \eta_m$, then $\Delta t(\mathbf{x}_0 \leftrightarrow \mathbf{x}_1)$ can be expressed as

$$\Delta t(\mathbf{x}_0 \leftrightarrow \mathbf{x}_1) = \frac{\eta_m}{c} \|\mathbf{x}_1 - \mathbf{x}_0\|. \quad (5)$$

The second form of delay occurs in the scattering events, and might occur from different sources, including electromagnetic phase shift, fluorescence and phosphorescence, or multiple scattering within the surface (or particle) microgeometry. To account for these sources of scattering delays, we introduce a temporal variable in the phase function as $f_s(\mathbf{x}, \vec{\omega}_i, \vec{\omega}, t)$, where t is the instant of light interacting with the particle before it is scattered. With those delays in place, we reformulate the RTE (Equations 1 and 2) introducing the temporal dependence as [Gla95]

$$L(\mathbf{x}, \vec{\omega}, t) = T_r(\mathbf{x}, \mathbf{x}_p) L_s(\mathbf{x}_p, \vec{\omega}, t - \Delta t_p) + \int_0^p \mu_s(\mathbf{x}_q) T_r(\mathbf{x}, \mathbf{x}_q) L_o(\mathbf{x}_q, \vec{\omega}, t - \Delta t_q) dq, \quad (6)$$

$$L_o(\mathbf{x}_q, \vec{\omega}, t) = \int_{-\infty}^t L_e(\mathbf{x}_q, \vec{\omega}, t') dt' + \int_{\mathcal{S}} \int_{-\infty}^t f_s(\mathbf{x}_q, \vec{\omega}_i, \vec{\omega}, t - t') L(\mathbf{x}_q, \vec{\omega}_i, t') dt' d\vec{\omega}_i, \quad (7)$$

with $\Delta t_p = \Delta t(\mathbf{x} \leftrightarrow \mathbf{x}_p)$ and $\Delta t_q = \Delta t(\mathbf{x} \leftrightarrow \mathbf{x}_q)$ (Equation (4)). L_s changes analogously. Note that we assume that the matter does not change at time-scales comparable to the speed of light, and therefore avoid any temporal dependence on μ_s and μ_t . Introducing temporal variation at such speeds would produce visible relativistic effects [WKR99, JMV*15].




Figure 2: (a) A photon emitted from the light source will take a time $t_{b_0} = \frac{\eta_m}{c}(s_1 + s_2 + s_3)$ to get to \mathbf{x}_b . (b) Radiance estimation in the medium is done by intersecting every ray against the photon beam map, and performing density estimations at the ray-beam intersections (red).

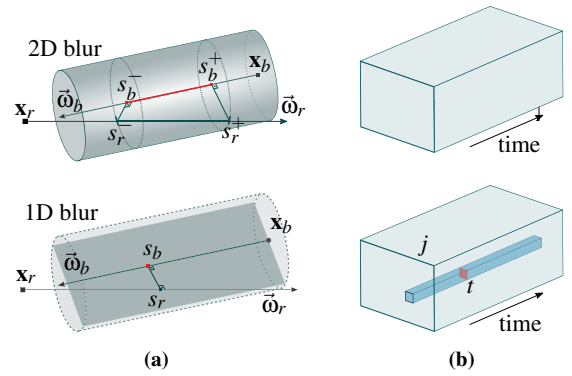


Figure 3: (a) Ray-beam intersection for density estimation using a 2D kernel (top) and 1D kernel (bottom). Time delays t_b, t_r within these spatial density estimations will depend on the ray-beam orientation the blur region intersections s_b, s_r , the speed of light, and the index of refraction of the media. (b) Radiance estimate of a single beam at pixel ij using a 2D blur generates a temporal footprint over a time interval $[t^-, t^+]$ (top) while radiance estimate using a 1D blur occurs at a single time instant t (bottom).

4. Transient Photon Beams

Photon beams [JNSJ11] provide a two-pass numerical solution for rendering participating media in steady state: In the first pass (Figure 2a), a series of random walk paths are traced from the light sources. These paths represent packages of light (photons) traveling through the medium. Every interaction of a photon within the medium is stored on a map as a *beam* with a direction $\vec{\omega}_b$, position \mathbf{x}_b and power Φ_b . In the second pass (Figure 2b), rays are traced from the camera against the scene, and Equation (1) is approximated by summing up the contribution of all near photon beams R_b of the eye ray defined by $r = (\mathbf{x}_r, -\vec{\omega}_r)$

$$L(\mathbf{x}_r, \vec{\omega}_r) \approx \sum_{b \in R_b} L_b(\mathbf{x}_r, \vec{\omega}_r), \quad (8)$$

where $L_b(\mathbf{x}_r, \vec{\omega}_r)$ is the contribution of photon beam b . Every photon beam b is considered to have certain radius R_b , and radiance

seen by a camera ray is computed by performing a density estimation on every ray-beam intersection. For 1D and 2D kernels, this radiance is computed as

$$L_b^{1D}(\mathbf{x}_r, \vec{\omega}_c) = K_{1D}(R_b) \Phi_b f_s(\theta_b) \mu_s \frac{e^{-\mu_s s_b} e^{-\mu_s s_r}}{\sin \theta_b}, \quad (9)$$

$$L_b^{2D}(\mathbf{x}_r, \vec{\omega}_r) = K_{2D}(R_b) \Phi_b f_s(\theta_b) \mu_s \frac{e^{-\mu_t (s_r^- - s_r^+)} (|\cos \theta_b| - 1) - 1}{e^{\mu_t (s_r^- + s_b^-)} \mu_t (|\cos \theta_b| - 1)}, \quad (10)$$

where the beam is defined by $\mathbf{x}_b + s_b \vec{\omega}_b$ and the ray is defined by $\mathbf{x}_r + s_r \vec{\omega}_r$ (see setups in Figure 3a).

4.1. Our algorithm

To generalize photon beams to the transient domain, we need to account for the duration of light paths. This requires considering propagation and scattering delays along the camera and light subpaths, but also the effect of time in the density estimation connecting these two subpaths.

Creating the photon map We compute the photon propagation as a standard random walk through the scene, which can be modeled using the subpath formulation defined by Jarabo et al. [JMM*14]. Let us define a light subpath $\bar{\mathbf{x}}_l = \mathbf{x}_0 \dots \mathbf{x}_{k-1}$, with k vertices, where \mathbf{x}_0 is the light source. This light path defines $k-1$ photon beams, in which a beam b_j is defined by its origin at $\mathbf{x}_{b_j} = \mathbf{x}_j$ and direction $\vec{\omega}_{b_j} = \frac{\mathbf{x}_{j+1} - \mathbf{x}_j}{\|\mathbf{x}_{j+1} - \mathbf{x}_j\|}$. Using Jarabo's definition of the path integral (and therefore of the contribution of the subpaths), we compute the flux of each photon as:

$$\Phi_{b_j} = \frac{f(\bar{\mathbf{x}}_j, \bar{\tau}_j)}{M p(\bar{\mathbf{x}}_j, \bar{\tau}_j)} = \frac{L_e(\mathbf{x}_0 \rightarrow \mathbf{x}_1, \tau_0) T(\bar{\mathbf{x}}_j, \bar{\tau}_j)}{M \prod_{i=0}^j p(\mathbf{x}_i, \tau_i)}, \quad (11)$$

with $\bar{\mathbf{x}}_j$ the subpath of $\bar{\mathbf{x}}_l$ up the vertex j , f the subpath contribution function, $\bar{\tau}_j = \tau_0 \dots \tau_j$ the sequence of time delays up to vertex j , M the number of photon random walks sampled, $L_e(\mathbf{x}_0 \rightarrow \mathbf{x}_1, \tau_0)$ the emission function, $p(\mathbf{x}_i, \tau_i)$ the probability density of sampling vertex \mathbf{x}_i with time delay τ_i . The throughput, $T(\bar{\mathbf{x}}_j, \bar{\tau}_j)$, of subpath (\mathbf{x}_i, τ_j) is defined as:

$$T(\bar{\mathbf{x}}_j, \bar{\tau}_j) = \left[\prod_{i=1}^{j-1} f_s(\mathbf{x}_i, \tau_j) \right] \left[\prod_{i=0}^{j-1} G(\mathbf{x}_i, \mathbf{x}_{i+1}) V(\mathbf{x}_i, \mathbf{x}_{i+1}) \right], \quad (12)$$

with $f_s(\mathbf{x}_i, \tau_j)$ the scattering event at vertex \mathbf{x}_i with delay τ_j , and $G(\mathbf{x}_i, \mathbf{x}_{i+1})$ and $V(\mathbf{x}_i, \mathbf{x}_{i+1})$ the geometry and visibility terms between vertices \mathbf{x}_i and \mathbf{x}_{i+1} , respectively. Finally, for transient state we need to know the instant t_{b_j} at which the photon beam is created (through emission or scattering), defined as:

$$t_{b_j} = \sum_{i=0}^{j-1} \tau_j + \sum_{i=0}^{j-1} \Delta t(\mathbf{x}_i, \mathbf{x}_{i+1}). \quad (13)$$

Rendering For rendering, we adapt Equation (8) to account for the temporal domain, as

$$L(\mathbf{x}_r, \vec{\omega}_r, t) \approx \sum_{b \in R_b} L_b(\mathbf{x}_r, \vec{\omega}_r, t), \quad (14)$$

with $L_b(\mathbf{x}_r, \vec{\omega}_r, t)$ the radiance estimation for beam b to ray t at instant t . In essence, $L_b(\mathbf{x}_r, \vec{\omega}_r, t)$ will return zero radiance if t is out of the temporal footprint of the density estimation kernel. Depending on the dimensionality of the density estimation, Jarosz and colleagues [JNSJ11] proposed three different estimators based on 3D, 2D and 1D kernels. Since the 3D kernel results impractical due to costly 3D convolutions, we focus on 1D and 2D kernels (Equations (9) and (10)), and extend them to transient state, assuming homogeneous media.

Kernel 2D We generalize Jarosz's et al.'s 2D estimate L_b^{2D} (Equation (10)) by introducing a temporal function $W(t)$ as

$$L_b^{2D}(\mathbf{x}_r, \vec{\omega}_r, t) = K_{2D}(R_b) \Phi_b f_s(\theta_b, t) \mu_s \frac{e^{-\mu_t (s_r^- - s_r^+)} (|\cos \theta_b| - 1) - 1}{e^{\mu_t (s_r^- + s_b^-)} \mu_t (|\cos \theta_b| - 1)} W_{2D}(t), \quad (15)$$

where $[s_r^-, s_r^+]$ are the limits of the ray-beam intersection (Figure 3a), θ_b is the angle between $\vec{\omega}_b$ and $\vec{\omega}_r$, and $K_{2D}(R_b)$ is a canonical 2D kernel with radius R_b . The temporal function $W_{2D}(t)$ models the temporal footprint of the 2D kernel as

$$W_{2D}(t) = \begin{cases} \frac{1}{t^+ - t^-} & \text{if } t \in (t^-, t^+) \\ 0 & \text{otherwise} \end{cases}, \quad (16)$$

where $t^- = t_b + t_r + \frac{\eta_m}{c} (s_r^- + s_b^-)$ and $t^+ = t_b + t_r + \frac{\eta_m}{c} (s_r^+ + s_b^+)$, and t_r and t_b are the initial times of the camera ray and beam, respectively. Note that due to transmittance, the photon energy varies as it travels across the blur region. Evenly distributing the integrated radiance L_b across this interval introduces temporal bias, in addition to the inherent spatial bias introduced by density estimation. However we observed this even distribution provides a good trade-off between bias, variance, and computational overhead.

Kernel 1D In the 1D kernel defined for density estimation by Jarosz et al. the spatial blur is performed over a line. Therefore, the energy of the beam is just spread on the ray on a single point at $r(s_r)$, from a single point of the beam $b(s_b)$ (see Figure 3a). In consequence, $s_r^\pm \rightarrow s_r$ and $s_b^\pm \rightarrow s_b$, which implies that $t^\pm \rightarrow t_{br}$, and the temporal function reduces to $W_{1D}(t - t_b) = \delta(t)$, with $\delta(t)$ the Dirac delta function. With that in place, we transform Jarosz et al. 1D estimate to

$$L_b^{1D}(\mathbf{x}_r, \vec{\omega}_r, t) = K_{1D}(R_b) \Phi_b f_s(\theta_b, t) \mu_s \frac{e^{-\mu_s s_b} e^{-\mu_s s_r}}{\sin \theta_b} \delta(t - t_b), \quad (17)$$

with $K_{1D}(R_b)$ a 1D kernel with radius R_b .

Implementation Since photon beams correspond to full photon trajectories, they allow us to estimate radiance at any position $\mathbf{x}_b + s \vec{\omega}_b$ of the beam, and therefore at any arbitrary time $t(\mathbf{x}_b + s \vec{\omega}_b)$. As mentioned, one-dimensional radiance estimate corresponds to a single time across the beam. In a traditional rendering process where camera rays are traced through view-plane pixels against the beams map, the temporal definition *within* a pixel will be proportional to the amount of samples per pixel taken. Additionally, 2D blur requires distributing every radiance estimate along a time interval, which reduces variance in the time dimension of a pixel at the expense of introducing additional temporal bias.

Finally, note that the temporal footprint of the density estimation might be arbitrarily small, so the probability of finding a beam b at an specific time might be very low. We alleviate this issue using path reuse via density estimation [JMM*14]. In particular, for the non-progressive results we use histogram temporal density estimation. In this technique, the samples in the temporal domain are reused across all frames by evaluating their contribution functions, which correspond to the temporal window covered by each frame.. In Section 5 we introduce temporal kernel-based density estimation, and combine it with the spatial density estimation of the beam.

5. Progressive Transient Photon Beams

By means of Equations (15) and (17) we have introduced temporal dependence on the spatial density estimations that use 2D and 1D kernels, respectively. These density estimations reduce variance at the expense of introducing bias in the results, which means both Equations (8) and (14) will not converge to the correct solution, even with an infinite number of photons M . To avoid this, progressive density estimation aims to provide a biased, yet consistent technique, that in the limit converges to the expected value (in other words, the bias vanishes in the limit). The key idea is to average several render passes with a finite number of photon random walks M , progressively reducing the bias in each iteration while allowing variance to slightly increase.

In order to fully leverage a progressive approach, we propose to combine our time-resolved spatial density estimations (Section 4) with additional *temporal* density estimations. While our time-resolved 2D spatial kernel implicitly performs a temporal blur over the interval $[t^-, t^+]$, it is coupled with the spatial blur. This does not allow to choose its own initial kernel size for the temporal density estimation, which is a desirable degree of freedom since the temporal resolution may not be proportional to the spatial one. In contrast, our time-resolved 1D spatial kernel does not perform a temporal blur, since the footprint is a single instant in time. As we show in the remainder of this section, this allows us to perform additional progressive temporal density estimations with an independent initial kernel size, while keeping the same two-dimensionality (1D spatial and 1D temporal). In the following, we introduce our spatio-temporal beam density estimation based on our time-resolved 1D kernel, and then present our progressive approach.

Spatio-Temporal Beam Estimation Jarabo et al. [JMM*14] showed that progressive density estimations in the temporal domain can in fact improve the convergence rate for transient rendering, in particular when compared with the histogram method used in Section 4 for rendering the temporal domain. To combine such approach with the (progressive) spatial density estimation in photon beams [JNT*11], we reformulate the 1D kernel in Equation (17), by convolving it with a 1D temporal kernel $K_{\mathcal{T}}(t)$ so that

$$L_b^{1D}(\mathbf{x}_r, \vec{\omega}_r, t) = K_{1D}(R_b) \Phi_b f_s(\theta_b, t) \mu_s \frac{e^{-\mu_s s_b} e^{-\mu_s s_r}}{\sin \theta_b} K_{\mathcal{T}}(t - t_b). \quad (18)$$

Progressive Transient Photon Beams We generalize the computation of $L(\mathbf{x}_r, \vec{\omega}_r, t)$ (Equation (14)) using an iterative estimator,

Algorithm 1 Pseudo-code of our progressive spatio-temporal density estimation.

```

 $L_n \leftarrow 0$ 
 $R_b \leftarrow R_0$ 
 $\mathcal{T} \leftarrow \mathcal{T}_0$ 
for  $i \in [0..N)$  do
     $r \leftarrow \text{traceRay}()$ 
     $B \leftarrow \text{beamsMap}()$  (Eqs. (6), (7), (11)-(13))
     $R_b \leftarrow R_b \sqrt{\frac{i+2/3}{i+1}}$  (Eq. (20), left)
     $\mathcal{T} \leftarrow \mathcal{T} \sqrt{\frac{i+2/3}{i+1}}$  (Eq. (20), right)
     $L_b \leftarrow 0$ 
    for  $b \in B$  do
         $L_b \leftarrow L_b + \text{radiance}(r, b, R_b, \mathcal{T})$  (Eq. (18))
    end for
     $L_n \leftarrow L_n + L_b$ 
end for

```

defined as

$$L(\mathbf{x}_r, \vec{\omega}_r, t) \approx \widehat{L}_n(\mathbf{x}_r, \vec{\omega}_r, t) = \frac{1}{n} \sum_{i=0}^n \sum_{b \in B_i} L_b(\mathbf{x}_r, \vec{\omega}_r, t) \quad (19)$$

with \widehat{L}_n the estimate of L after n iterations, and B_i the set of photon beams in iteration i . Note that the previous equation assumes that the camera ray r is the same for all iterations. That is not necessarily true (and in fact it is not) but for simplicity we express this way.

The error of the estimate \widehat{L}_n is defined by its bias and variance, which as shown in Appendix B is dependent on the bandwidth of the spatial and temporal kernels. In particular, the variance of the error increases linearly with the bandwidth of the kernels, while bias is reduced at the same rate. Then, on each iteration we reduce the bias by allowing the variance to increase at a controlled rate of $(i+1)/(i+\alpha)$, with $\alpha \in [0, 1]$ being a parameter that controls how much the variance is allowed to increase at each iteration. To achieve that reduction, on each iteration $i+1$ we reduce the footprint of kernels K_{1D} and $K_{\mathcal{T}}$ ($R_{b|i}$ and \mathcal{T}_i) by

$$\frac{R_{b|i+1}}{R_{b|i}} = \left(\frac{i+\alpha}{i+1} \right)^{\beta_R}, \quad \frac{\mathcal{T}_{i+1}}{\mathcal{T}_i} = \left(\frac{i+\alpha}{i+1} \right)^{\beta_{\mathcal{T}}}, \quad (20)$$

where β_R and $\beta_{\mathcal{T}}$ control the individual reduction ratio of each kernel, with $\beta_{\mathcal{T}} = 1 - \beta_R$. A pseudo code of the main steps of our progressive approach can be found in Algorithm 1. In the following, we analyze the convergence rate of the method, and compute the optimal values for the parameters α , $\beta_{\mathcal{T}}$ and β_R .

Convergence analysis We analyze the convergence of the algorithm as a function of the *asymptotic mean squared error* (AMSE) defined as

$$\text{AMSE}(\widehat{L}_n) = \text{Var}[\widehat{L}_n] + \text{E}[\varepsilon_n]^2, \quad (21)$$

where $\text{Var}[\widehat{L}_n]$ is the variance of the estimate and $\text{E}[\varepsilon_n]$ is the bias at iteration n (see Appendix A). As shown in Appendix C, the variance converges with rate

$$\text{Var}[\widehat{L}_n] \approx O(n^{-1}) + O(n^{-\alpha}) = O(n^{-\alpha}), \quad (22)$$




Figure 4: Steady-state renders for the scenes CORNELL SPHERES (Figure 5), MIRRORS (Figure 6), and PUMPKIN (Figure 7).

while the bias converges with rate

$$E[\epsilon_n] = O(n^{1-\alpha})^{-2\beta\tau} + O(n^{1-\alpha})^{2\beta\tau-2}. \quad (23)$$

Plugging Equation (22) and (23) into Equation (21), we can model the AMSE as

$$\text{AMSE}(\hat{L}_n) = O(n^{-\alpha}) + \left(O(n^{1-\alpha})^{-2\beta\tau} + O(n^{1-\alpha})^{2\beta\tau-2}\right)^2 \quad (24)$$

Finally, by minimizing Equation (24) (see Appendix D) we obtain the values for optimal asymptotic convergence $\beta\tau = 1/2$ and $\alpha = 2/3$, which by substitution gives us the final asymptotic convergence rate of our progressive transient photon beams

$$\text{AMSE}(\hat{L}_n) = O(n^{-\frac{2}{3}}). \quad (25)$$

6. Results

In the following we illustrate the results of our proposed method in five scenes: CORNELL SPHERES, MIRRORS, PUMPKIN, SOCCER [SZLG10], PUMPKIN, and JUICE. See Figures 1 (right), 4, and 8 (left) for steady-state renders of the scenes. Results of Figures 5 and 6 were taken on a desktop PC with Intel i7 and 4GB RAM using a transient 2D kernel (Equation 15). Figures 1, 7, and 8 were rendered on an Intel Xeon E5 with 256GB RAM, using our progressive spatio-temporal kernel density estimations (Section 5) derived from the transient spatial 1D kernel (Equation 17). In each iteration, we use a fixed radius for our spatio-temporal density estimators (instead of using a nearest neighbor approach). Please refer to the supplemental video for the full sequences of all the scenes.

Figure 5 shows a Cornell box filled with a scattering medium, and demonstrates the effect of *camera unwarping* [VWJ*13] when rendering. Camera unwarping is an intuitive way of visualizing how light propagates *locally* on the scene without accounting for the time light takes to reach the camera. The scene consists of a diffuse Cornell box with a point light on the top, a glass refractive sphere (top, IOR = 1.5) and a mirror sphere (bottom). While Figure 5b shows the real propagation of light—including camera time—, Figure 5a depicts more intuitively how light comes out from the point light, travels through the refractive sphere, and the generated caustic bounces on the mirror sphere. Note how in the top sequence we can clearly see how light is slowed down through the glass sphere due to the higher index of refraction. We can also observe multiple scattered light (particularly noticeable in frames $t=4\text{ns}$ and $t=6\text{ns}$) as a secondary wavefront.

Figure 6 compares visualizations of light propagation within the




Figure 5: Comparison of CORNELL SPHERES scene using camera-unwarping (top), where we do not take into account the camera time, and real propagation of light (bottom). In the bottom row the shape of the wavefront is altered by the camera time, as if we were scanning the scene from the viewpoint towards the furthest parts of the scene. Camera unwarping on the other hand illustrates more intuitively how light propagates locally.

MIRRORS scene using Heaviside and Dirac delta light emission. The scene is composed by two colored mirrors and a glass sphere with IOR = 1.5, and was rendered using the previously mentioned camera unwarping. We can observe how delta emission generates wavefronts that go through the ball and bounce in the mirrors, creating wavefront holes where constant emission creates medium shadows. In the last frame of the top row Delta emission clearly depicts the slowed down caustic through the glass ball respect to the main wavefront.

Our progressive method combines time-resolved 1D spatial kernels of photon beams and temporal density estimations, reducing bias while providing consistent solutions in the limit with an optimal convergence rate of $O(n^{-\frac{2}{3}})$. In Figure 7 we analyze its convergence with respect to progressive transient path tracing with temporal KDE [JMM*14] (PTPT). In the middle graph we show the temporal profile on a single pixel for both our algorithm and PTPT after 4096 equal-time iterations, where both algorithms converge to the reference solution taken with transient path tracing (no temporal KDE) with 64 million samples. While PTPT presents faster convergence (see Figure 7, right graph), our algorithm presents a better behavior over time where variance increases due to the lack of samples (center graph). Additionally, it requires much fewer iterations than PTPT to achieve a similar MSE (see log-log right graph).

In Figure 1 we show a more complex scenario, with different caustics rendered, with our progressive algorithm. It contains a smooth dielectric figurine with different transmission albedos placed within a participating medium with an isotropic phase function. Our method is capable of handling complex caustics transmitted from light sources through the player, and then through the ball. Our algorithm progressively reduces bias and variance to provide a consistent solution.

Finally in Figure 8 we illustrate a setup combining different media properties, and specular refractive and reflective materials. The liquid has a very forward phase function, making the light first travel through the direction of the stream ($t = 4.6$ ns), and then




Figure 6: Comparison between Dirac delta (top) and continuous emission (bottom). Dirac delta emission lets us see how a pulse of light travels and scatters across the scene, depicting the light wavefronts bouncing on the mirrors and going through the glass ball. Continuous emission shows how light is emitted until it reaches every point in the scene, as if we were taking a picture with a camera at very slow-motion.




Figure 7: The PUMPKIN scene shows a jack o’lantern embedding a point light that creates hard shadows through the holes. The left frames show a sequence of the time-resolved renders after 4096 iterations of our algorithm (10k beams / iteration), and temporal KDE on a progressive transient path tracer (PTPT, 16spp / iteration) [JMM*14]. The middle plot compares the whole temporal footprint at the pink marker. Reference solution (dark grey) was obtained with a transient path tracer (no KDE) using 64M samples per pixel. Right plot shows MSE convergence with respect to the number of progressive iterations (in log-log scale), at 1 minute/iteration on each algorithm. As expected, the convergence of our method ($O(n^{-\frac{2}{3}})$) is slower than PTPT ($O(n^{-\frac{4}{5}})$); however, as shown in the equal-time comparison, our algorithm presents better temporal behavior with much less variance on later timings.




Figure 8: We illustrate the potential of our method in the JUICE scene [Bit16a], which presents a scene very difficult to render for path tracing methods, but well-handled by photon-based methods. The scene is filled by a thin participating medium, while the glass contains ruby grapefruit juice as measured by Narasimhan et al. [NGD*06]. The highly forward phase function of the juice, as well as the delta interactions on the glass, ice cubes, and the mirror floor surface, generate complex caustic patterns which our method is able to simulate in transient state. Bottom row has increased exposure respect to top row to show the radiance at later timings.

going through the liquid inside the glass ($t = 5.1\text{ns}$ to $t = 6.3\text{ns}$). The mirror surface makes the light to bounce back to the surrounding medium as a caustic through the water spills and ice cubes at $t = 5.1\text{ns}$ and $t = 6.6\text{ns}$. Note that these are not fully observable in the steady-state render (left) due to the accumulated radiance from the surrounding medium and the adjusted exposure of the image.

7. Conclusions

In this paper we have presented a robust progressive method for efficiently rendering transient light transport with consistent re-

sults. We derived our method based on progressive photon beams [JNT*11], extending its density estimators to account for light time-of-flight, and deriving a new progressive scheme. We then compute the convergence of the method, and derive the parameters for optimal asymptotic convergence. Our results demonstrate that combining continuous photon trajectories in transient state and our optimal spatio-temporal convergence rates allow to robustly compute a noise-free solutions to the time-resolved RTE for complex light paths. We believe that our work might be very useful for developing new techniques for transient imaging and reconstruction

in media, as well as to obtain new insights on time-resolved light transport.

As future work it would be interesting to analyze more thoroughly the optimal performance and kernels for variance reduction and bias impact in transient state, under varying media characteristics. In addition, extending our method to leverage recent advances in media transport, such as transient-state adaptations of higher-dimensional photon estimators [BJ17] as well as hybrid techniques [KGH*14], could improve performance of time-resolved rendering for a general set of geometries and media characteristics.

Acknowledgments

We want to thank the reviewers for their insightful comments. This project has been funded by DARPA (project REVEAL), the European Research Council (ERC) under the EU's Horizon 2020 research and innovation programme (project CHAMELEON, grant No 682080), the Spanish Ministry of Economy and Competitiveness (project TIN2016-78753-P), the BBVA Foundation, and the Gobierno de Aragón.

Appendix A: Error in Transient Progressive Photon Beams

Here we analyze the consistency of the transient progressive photon beams algorithm described in Section 5. For our analysis on the error of the estimate, we use the *asymptotic mean squared error* (AMSE) defined as

$$\text{AMSE}(\widehat{L}_n) = \text{Var}[\widehat{L}_n] + \text{E}[\varepsilon_n]^2, \quad (26)$$

where $\text{Var}[\widehat{L}_n]$ is the variance of the estimate and $\text{E}[\varepsilon_n]$ is the bias at iteration n . We model $\text{Var}[\widehat{L}_n]$ as [KZ11]

$$\text{Var}[\widehat{L}_n] = \frac{1}{n} \text{Var}[\Psi L] + \frac{1}{n^2} \sum_{j=1}^n \text{Var}[\Psi \varepsilon_j], \quad (27)$$

where Ψ is the contribution of the eye ray, and ε_j is the bias for iteration j . The first term is the standard variance of the Monte Carlo estimate, which is unaffected by the kernel. The second term, on the other hand, is the variance of the error, and is dependent on density estimation. On the other hand, the estimated value of the error (bias) $\text{E}[\widehat{L}_n]$ is defined as

$$\text{E}[\widehat{L}_n] = L + \text{E}[\Psi] \text{E}[\varepsilon_n], \quad (28)$$

where $\text{E}[\varepsilon_n]$ is the bias of the estimator after n steps:

$$\text{E}[\varepsilon_n] = \frac{1}{n} \sum_{j=1}^n \text{E}[\varepsilon_j], \quad (29)$$

with $\text{E}[\varepsilon_j]$ the expected error at iteration j . In the following, we first derive the variance and expected value of the error for a single iteration. Then, we analyze the asymptotic behavior of these terms, and compute the values for optimal convergence for $\beta_{\mathcal{T}}$, β_R and α .

Appendix B: Variance and Expected Value of the Error of the Time-Resolved Beam Radiance Estimate

We first analyze the variance and expected value of the error (bias) introduced by the radiance estimate at each iteration. Let us first define the error in each iteration as:

$$\begin{aligned} \varepsilon &= \widehat{L}_n(\mathbf{x}_r, \vec{\omega}_r, t) - L(\mathbf{x}_r, \vec{\omega}_r, t) \\ &= \sum_{i=1}^M K_{1D}(R_b) K_{\mathcal{T}}(t - t_i) \Phi_i - L(\mathbf{x}_r, \vec{\omega}_r, t). \end{aligned} \quad (30)$$

Variance We first define the variance of the error $\text{Var}[\varepsilon]$ as (in the following, we omit dependences for clarity):

$$\begin{aligned} \text{Var}[\varepsilon] &= \text{Var}\left[\sum_{i=1}^M K_{1D} K_{\mathcal{T}} \Phi - L\right] \\ &= (\text{Var}[K_{1D}] + \text{E}[K_{1D}]^2) (\text{Var}[K_{\mathcal{T}}] + \text{E}[K_{\mathcal{T}}]^2) \\ &\quad (\text{Var}[\Phi] + \text{E}[\Phi]^2) - \text{E}[K_{1D}]^2 \text{E}[K_{\mathcal{T}}]^2 \text{E}[\Phi]^2, \end{aligned} \quad (31)$$

In order to compute the variance of the error $\text{Var}[\varepsilon]$ we need to make a set of assumptions: First, we assume that the beams' probability density is constant within the kernel K_{1D} in the spatial domain [JNT*11], and within $K_{\mathcal{T}}$ in the temporal domain [JMM*14]. We denote these probabilities as p_{R_b} and $p_{\mathcal{T}}$ respectively. We also assume that the distance between view ray and photon beam, time t_b and beams' energy Φ_i are independent samples of the random variables D , T and Φ , respectively, which are mutually independent. Finally, we assume that D and T have probability densities p_{R_b} and $p_{\mathcal{T}}$.

With these assumptions, and taking into account that $\text{E}[K_{1D}] = p_{R_b}$ and $\text{E}[K_{\mathcal{T}}] = p_{\mathcal{T}}$, we can model the variance introduced by the temporal kernel $\text{Var}[K_{\mathcal{T}}]$ as [JMM*14]

$$\text{Var}[K_{\mathcal{T}}] = \frac{p_{\mathcal{T}}}{T} \int_{\mathbb{R}} k_{\mathcal{T}}(\psi)^2 d\psi - p_{\mathcal{T}}^2, \quad (32)$$

where we express $K_{\mathcal{T}}$ as a canonical kernel $k_{\mathcal{T}}$ with unit integral such that $K_{\mathcal{T}}(\xi) = k_{\mathcal{T}}(\xi/T) T^{-1}$. Analogously, $\text{Var}[K_{1D}]$ is [JNT*11]:

$$\text{Var}[K_{1D}] = \frac{p_{R_b}}{R_b} \int_{\mathbb{R}} k_{1D}(\psi)^2 d\psi - p_{R_b}^2. \quad (33)$$

This allows us to express the variance of the error $\text{Var}[\varepsilon]$ as:

$$\text{Var}[\varepsilon] \approx \left(\text{Var}[\Phi] + \text{E}[\Phi]^2 \right) \left(\frac{p_{R_b}}{R_b} \mathcal{C}_{1D} \right) \left(\frac{p_{\mathcal{T}}}{T} \mathcal{C}_{\mathcal{T}} \right), \quad (34)$$

where \mathcal{C}_{1D} and $\mathcal{C}_{\mathcal{T}}$ are kernel-dependent constants. The last term can be neglected by assuming that the kernels cover small areas in their respective domains, which effectively means that $\mathcal{C}_{1D} \gg p_{R_b}$ and $\mathcal{C}_{\mathcal{T}} \gg p_{\mathcal{T}}$. Equation (34) shows that for transient density estimation, the variance $\text{Var}[\varepsilon]$ is inversely proportional to $R_b T$.

Bias Bias at each iteration j is defined as the expected value of the error $\text{E}[\varepsilon_j]$ as

$$\begin{aligned} \text{E}[\varepsilon_j] &= \text{E}\left[\sum_{i=1}^M K_{1D} K_{\mathcal{T}} \Phi - L\right] \\ &= \text{E}[K_{1D}] \text{E}[K_{\mathcal{T}}] \text{E}[\Phi] - L. \end{aligned}$$

Using a second-order expansion of $p_{\mathcal{T}}$ and p_{R_b} , instead of the zeroth-order used when modeling variance, we can express the expected value of $K_{\mathcal{T}}$ as [JMM*14]

$$E[K_{\mathcal{T}}] \approx p_{\mathcal{T}} + \mathcal{T}^2 \int_{\mathbb{R}} k_{\mathcal{T}}(\psi) O(\|\psi\|^2) d\psi = p_{\mathcal{T}} + \mathcal{T}^2 C_{\mathcal{T}}^{ii}, \quad (35)$$

while the expected value of K_{1D} is [JNT*11]

$$E[K_{1D}] \approx p_{R_b} + R_b \int_{\mathbb{R}^2} k_{1D}(\psi) O(\|\psi\|^2) d\psi = p_{R_b} + R_b C_{1D}^{ii}, \quad (36)$$

where $C_{\mathcal{T}}^{ii}$ and C_{1D}^{ii} are constants dependent on the higher-order derivatives of the spatio-temporal light distribution. Using (35) and (36), and $L = p_{R_b} p_{\mathcal{T}} E[\Phi]$ we finally compute $E[\varepsilon_j]$ for iteration j as

$$\begin{aligned} E[\varepsilon_j] &\approx (p_{R_b} + R_b^2 C_{1D}^{ii})(p_{\mathcal{T}} + \mathcal{T}^2 C_{\mathcal{T}}^{ii}) E[\Phi] - p_{R_b} p_{\mathcal{T}} E[\Phi] \\ &= E[\Phi] (p_{R_b} \mathcal{T}^2 C_{\mathcal{T}}^{ii} + p_{\mathcal{T}} R_b^2 C_{1D}^{ii} + \mathcal{T}^2 C_{\mathcal{T}}^{ii} R_b^2 C_{1D}^{ii}). \end{aligned} \quad (37)$$

Appendix C: Convergence Analysis of Progressive Transient Photon Beams

Based on the expressions for $\text{Var}[\varepsilon]$ and $E[\varepsilon_j]$ defined above (Equations (34) and (37)), we can know derive the asymptotic behaviour of Equation (21). For that, we will compute the variance $\text{Var}[\widehat{L}_n]$ and bias $E[\varepsilon_n]$ after n iterations.

Variance Assuming that the random variables Ψ and ε_j are independent, we model the variance of the estimator $\text{Var}[\widehat{L}_n]$ in Equation (27) as [KZ11]:

$$\begin{aligned} \text{Var}[\widehat{L}_n] &= \frac{1}{n} \text{Var}[\Psi L] + \frac{1}{n^2} \sum_{j=1}^n \text{Var}[\Psi \varepsilon_j] \\ &= \frac{1}{n} \text{Var}[\Psi L] + \text{Var}[\Psi] \frac{1}{n^2} \sum_{j=1}^n \text{Var}[\varepsilon_j] + \\ &\quad E[\Psi]^2 \frac{1}{n^2} \sum_{j=1}^n \text{Var}[\varepsilon_j] + \text{Var}[\Psi] \frac{1}{n^2} \sum_{j=1}^n E[\varepsilon_j]^2. \end{aligned} \quad (38)$$

Following [KD13], we can approximate $\text{Var}[\varepsilon_n]$ as a function of the variance at the first iteration $\text{Var}[\varepsilon_1]$ as:

$$\text{Var}[\varepsilon_n] \approx \frac{\text{Var}[\varepsilon_1]}{(2 - \alpha)n^\alpha} = O(n^{-\alpha}). \quad (39)$$

Finally, by applying $\text{Var}[\varepsilon_n]$ and asymptotic simplifications, we can formulate $\text{Var}[\widehat{L}_n]$ (39) as:

$$\begin{aligned} \text{Var}[\widehat{L}_n] &\approx \frac{1}{n} \text{Var}[\Psi L] + E[\Psi]^2 \text{Var}[\varepsilon_n] \\ &\approx \frac{1}{n} \text{Var}[\Psi L] + \frac{\text{Var}[\varepsilon_1]}{(2 - \alpha)n^\alpha} \\ &= O(n^{-1}) + O(n^{-\alpha}) = O(n^{-\alpha}). \end{aligned} \quad (40)$$

Bias The expected value of the error $E[\varepsilon_n]$ is modeled in Equation (28) as a function of the averaged bias introduced at each iteration $E[\varepsilon_j]$ (37). Computing the kernels' bandwidth \mathcal{T}_j and R_{b_j} at

iteration j by expanding Equation (20) as a function of their initial value by we get

$$\mathcal{T}_j = \mathcal{T}_1 (j \alpha B(\alpha, j))^{-\beta_{\mathcal{T}}}, \quad (41)$$

$$R_{b_j} = R_{b1} (j \alpha B(\alpha, j))^{-\beta_{R_b}}, \quad (42)$$

where $B(x, y)$ is the Beta function. Using (41) and (42) in Equation (37) we can express $E[\varepsilon_j]$ as a function of the initial kernel bandwidths

$$\begin{aligned} E[\varepsilon_j] &= E[\Phi] p_{R_b} C_{\mathcal{T}}^{ii} \mathcal{T}_1^2 \Theta(j^{1-\alpha})^{-2\beta_{\mathcal{T}}} \\ &\quad + E[\Phi] p_{\mathcal{T}} C_{1D}^{ii} R_{b1}^2 \Theta(j^{1-\alpha})^{-2\beta_{R_b}} \\ &\quad + E[\Phi] C_{\mathcal{T}}^{ii} C_{1D}^{ii} \mathcal{T}_1^2 R_{b1}^2 \Theta(j^{1-\alpha})^{-2(\beta_{\mathcal{T}} + \beta_{R_b})}. \end{aligned} \quad (43)$$

Finally, we use $\sum_{j=1}^n \Theta(j^x) = n O(n^x)$ to plug Equation (43) into Equation (29) to get the asymptotic behavior of $E[\varepsilon_n]$ in transient progressive photon beams:

$$E[\varepsilon_n] = O(n^{1-\alpha})^{-2\beta_{\mathcal{T}}} + O(n^{1-\alpha})^{-2\beta_{R_b}} + O(n^{1-\alpha})^{-2(\beta_{\mathcal{T}} + \beta_{R_b})},$$

which, by using the equality $\beta_{R_b} = 1 - \beta_{\mathcal{T}}$, becomes:

$$\begin{aligned} E[\varepsilon_n] &= O(n^{1-\alpha})^{-2\beta_{\mathcal{T}}} + O(n^{1-\alpha})^{2\beta_{\mathcal{T}}-2} + O(n^{1-\alpha})^{-2} \\ &= O(n^{1-\alpha})^{-2\beta_{\mathcal{T}}} + O(n^{1-\alpha})^{2\beta_{\mathcal{T}}-2}. \end{aligned} \quad (44)$$

Appendix D: Minimizing Asymptotic Mean Squared Error

Using the asymptotic expression for variance and bias in Equations (40) and (44), we can express the AMSE (21) as

$$AMSE(\widehat{L}_n) = O(n^{-\alpha}) + \left(O(n^{1-\alpha})^{-2\beta_{\mathcal{T}}} + O(n^{1-\alpha})^{2\beta_{\mathcal{T}}-2} \right)^2. \quad (45)$$

which is a function of the parameters α and $\beta_{\mathcal{T}}$. Given that the variance is independent of $\beta_{\mathcal{T}}$, we first obtain the optimal value for this parameter that yields the highest convergence rate of the bias $E[\varepsilon_n]$. We differentiate Equation (44), apply asymptotic simplifications and equating to zero, we obtain the optimal value $\beta_{\mathcal{T}} = 1/2$. By plugging this value in Equation (45), we obtain:

$$AMSE(\widehat{L}_n) = O(n^{-\alpha}) + O(n^{-2(1-\alpha)}). \quad (46)$$

Finally, by finding the minimum again with respect to α we get the optimal parameter $\alpha = 2/3$, which results in the optimal convergence rate of the AMSE for our transient progressive photon beams as

$$AMSE(\widehat{L}_n) = O(n^{-\frac{2}{3}}) + O(n^{-2(1-\frac{2}{3})}) = O(n^{-\frac{2}{3}}). \quad (47)$$

References

- [ABW14] AMENT M., BERGMANN C., WEISKOPF D.: Refractive radiative transfer equation. *ACM Trans. Graph.* 33, 2 (2014), 17:1–17:22. 2, 3
- [AGJ17] ARELLANO V., GUTIERREZ D., JARABO A.: Fast back-projection for non-line of sight reconstruction. *Optics Express* 25, 10 (2017), 11574–11583. 1
- [BG70] BELL G. I., GLASSTONE S.: *Nuclear reactor theory*. Tech. rep., US Atomic Energy Commission, Washington, DC (United States), 1970. 2

- [BH04] BUSCK J., HEISELBERG H.: Gated viewing and high-accuracy three-dimensional laser radar. *Applied Optics* 43, 24 (2004), 4705–4710. 1
- [Bit16a] BITTERLI B.: Rendering resources, 2016. <https://benedikt-bitterli.me/resources/>. 7
- [Bit16b] BITTERLI B.: Virtual femto photography. <https://benedikt-bitterli.me/femto.html>, 2016. 2
- [BJ17] BITTERLI B., JAROSZ W.: Beyond points and beams: Higher-dimensional photon samples for volumetric light transport. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 36, 4 (2017). 2, 8
- [Bus05] BUSCK J.: Underwater 3-d optical imaging with a gated viewing laser radar. *Optical Engineering* 44, 11 (2005), 44 – 44 – 7. 1
- [Cha60] CHANDRASEKHAR S.: *Radiative Transfer*. Dover, 1960. 2
- [CJ02] CAMMARANO M., JENSEN H. W.: Time dependent photon mapping. In *Eurographics Workshop on Rendering* (2002). 2
- [CPH53] CASE K. M., PLACZEK G., HOFFMANN F.: Introduction to the theory of neutron diffusion, v. 1. 2
- [DM79] DUDERSTADT J. J., MARTIN W. R.: Transport theory. *Transport theory* (1979). 2
- [GKDS12] GEORGIEV I., KRIVÁNEK J., DAVIDOVIČ T., SLUSALLEK P.: Light transport simulation with vertex connection and merging. *ACM Trans. Graph.* 31, 6 (2012), 192:1–192:10. 2
- [Gla95] GLASSNER A. S.: *Principles of digital image synthesis*, vol. 1. Elsevier, 1995. 2, 3
- [GMAS05] GUTIERREZ D., MUÑOZ A., ANSON O., SERON F.: Non-linear volume photon mapping. In *Eurographics Symposium on Rendering* (2005). 3
- [GRv*16] GRUSON A., RIBARDIÈRE M., ŠIK M., VORBA J., COZOT R., BOUATOUCH K., KRIVÁNEK J.: A spatial target function for metropolis photon tracing. *ACM Trans. Graph.* 36, 1 (2016). 2
- [HJG*13] HACHISUKA T., JAROSZ W., GEORGIEV I., KAPLANYAN A., NOWROUZEZHRAI D.: State of the art in photon density estimation. In *ACM SIGGRAPH ASIA 2013 Courses* (2013). 2
- [HOJ08] HACHISUKA T., OGAKI S., JENSEN H. W.: Progressive photon mapping. *ACM Trans. Graph. (TOG)* 27, 5 (2008), 130. 2
- [HPJ12] HACHISUKA T., PANTALEONI J., JENSEN H. W.: A path space extension for robust light transport simulation. *ACM Trans. Graph. (TOG)* 31, 6 (2012), 191. 2
- [JA18] JARABO A., ARELLANO V.: Bidirectional rendering of vector light transport. *Computer Graphics Forum* 37, 6 (2018), 96–105. 2
- [Jar12] JARABO A.: *Femto-photography: Visualizing light in motion*. Master's thesis, Universidad de Zaragoza, 2012. 2
- [JC98] JENSEN H. W., CHRISTENSEN P. H.: Efficient simulation of light transport in scenes with participating media using photon maps. In *SIGGRAPH '98* (1998), ACM, pp. 311–320. 2
- [Jen01] JENSEN H. W.: *Realistic Image Synthesis Using Photon Mapping*. AK Peters, 2001. 2
- [JMM*14] JARABO A., MARCO J., MUÑOZ A., BUISAN R., JAROSZ W., GUTIERREZ D.: A framework for transient rendering. *ACM Trans. Graph.* 33, 6 (2014), 177:1–177:10. 2, 4, 5, 6, 7, 8, 9
- [JMMG17] JARABO A., MASIA B., MARCO J., GUTIERREZ D.: Recent advances in transient imaging: A computer graphics and vision perspective. *Visual Informatics* 1, 1 (2017), 65–79. 1
- [JMV*15] JARABO A., MASIA B., VELTEN A., BARSİ C., RASKAR R., GUTIERREZ D.: Relativistic effects for time-resolved light transport. *Computer Graphics Forum* 34, 8 (2015), 1–12. 3
- [JNSJ11] JAROSZ W., NOWROUZEZHRAI D., SADEGHI I., JENSEN H. W.: A comprehensive theory of volumetric radiance estimation using photon points and beams. *ACM Trans. Graph.* 30, 1 (Feb. 2011), 5:1–5:19. 2, 3, 4
- [JNT*11] JAROSZ W., NOWROUZEZHRAI D., THOMAS R., SLOAN P.-P., ZWICKER M.: Progressive photon beams. *ACM Trans. Graph. (TOG)* 30, 6 (2011), 181. 2, 5, 7, 8, 9
- [JRJ11] JAKOB W., REGG C., JAROSZ W.: Progressive expectation–maximization for hierarchical volumetric photon mapping. *Computer Graphics Forum (Proc. of Eurographics Symposium on Rendering)* 30, 4 (jun 2011). 2
- [JZJ08] JAROSZ W., ZWICKER M., JENSEN H. W.: The beam radiance estimate for volumetric photon mapping. In *ACM SIGGRAPH 2008 classes* (2008), ACM, p. 3. 2
- [Kaj86] KAJIYA J. T.: The rendering equation. In *SIGGRAPH* (1986). 3
- [KD13] KAPLANYAN A. S., DACHSBACHER C.: Adaptive progressive photon mapping. *ACM Trans. Graph. (TOG)* 32, 2 (2013), 16. 2, 9
- [KGH*14] KRIVÁNEK J., GEORGIEV I., HACHISUKA T., VÉVODA P., ŠIK M., NOWROUZEZHRAI D., JAROSZ W.: Unifying points, beams, and paths in volumetric light transport simulation. *ACM Trans. Graph. (Proceedings of SIGGRAPH)* 33, 4 (July 2014). 2, 8
- [KPM*16] KLEIN J., PETERS C., MARTÍN J., LAURENZIS M., HULLIN M. B.: Tracking objects outside the line of sight using 2D intensity images. *Scientific Reports* 6 (2016). 2
- [KZ11] KNAUS C., ZWICKER M.: Progressive photon mapping: A probabilistic approach. *ACM Trans. Graph. (TOG)* 30, 3 (2011), 25. 2, 8, 9
- [Mar13] MARCO J.: *Transient Light Transport in Participating Media*. Master's thesis, Universidad de Zaragoza, 2013. 2
- [MHM*17] MARCO J., HERNANDEZ Q., MUÑOZ A., DONG Y., JARABO A., KIM M., TONG X., GUTIERREZ D.: DeepToF: Off-the-shelf real-time correction of multipath interference in time-of-flight imaging. *ACM Transactions on Graphics (SIGGRAPH Asia 2017)* 36, 6 (2017). 1
- [MJGJ17] MARCO J., JAROSZ W., GUTIERREZ D., JARABO A.: Transient photon beams. *Spanish Computer Graphics Conference (CEIG)* (2017). 2
- [MNJK13] MEISTER S., NAIR R., JÄHNE B., KONDERMANN D.: *Photon Mapping based Simulation of Multi-Path Reflection Artifacts in Time-of-Flight Sensors*. Tech. rep., Heidelberg Collaboratory for Image Processing, 2013. 2
- [NGD*06] NARASIMHAN S. G., GUPTA M., DONNER C., RAMAMOORTHY R., NAYAR S. K., JENSEN H. W.: Acquiring scattering properties of participating media by dilution. In *ACM Transactions on Graphics (TOG)* (2006), vol. 25, ACM, pp. 1003–1012. 7
- [NGHJ18] NOVÁK J., GEORGIEV I., HANIKA J., JAROSZ W.: Monte carlo methods for volumetric light transport simulation. *Computer Graphics Forum (Proceedings of Eurographics - State of the Art Reports)* 37, 2 (May 2018). 2
- [NZV*11] NAIK N., ZHAO S., VELTEN A., RASKAR R., BALA K.: Single view reflectance capture using multiplexed scattering and time-of-flight imaging. *ACM Trans. Graph.* 30 (2011), 171:1–171:10. 1
- [OHX*14] O'TOOLE M., HEIDE F., XIAO L., HULLIN M. B., HEIDRICH W., KUTULAKOS K. N.: Temporal frequency probing for 5D transient analysis of global light transport. *ACM Trans. Graph.* 33, 4 (2014), 87:1–87:11. 2
- [PBSC14] PITTS P., BENEDETTI A., SLANEY M., CHOU P.: *Time of Flight Tracer*. Tech. rep., Microsoft, 2014. 2
- [SJ09] SPENCER B., JONES M. W.: Into the blue: Better caustics through photon relaxation. *Computer Graphics Forum* 28, 2 (2009), 319–328. 2
- [SSD08] SMITH A., SKORUPSKI J., DAVIS J.: *Transient Rendering*. Tech. Rep. UCSC-SOE-08-26, School of Engineering, University of California, Santa Cruz, 2008. 2
- [SZLG10] SUN X., ZHOU K., LIN S., GUO B.: Line space gathering for single scattering in large scenes. In *ACM Transactions on Graphics (TOG)* (2010), vol. 29, ACM, p. 54. 6

- [VWG*12] VELTEN A., WILLWACHER T., GUPTA O., VEERARAGHAVAN A., BAWENDI M. G., RASKAR R.: Recovering three-dimensional shape around a corner using ultrafast time-of-flight imaging. *Nature Communications*, 3 (2012). 1
- [VWJ*13] VELTEN A., WU D., JARABO A., MASIA B., BARSÌ C., JOSHI C., LAWSON E., BAWENDI M., GUTIERREZ D., RASKAR R.: Femto-photography: Capturing and visualizing the propagation of light. *ACM Trans. Graph.* 32, 4 (2013), 44:1–44:8. 1, 6
- [Wil71] WILLIAMS M. M. R.: Mathematical methods in particle transport theory. 2
- [WJS*18] WU R., JARABO A., SUO J., DAI F., ZHANG Y., DAI Q., GUTIERREZ D.: Adaptive polarization-difference transient imaging for depth estimation in scattering media. *Opt. Lett.*, 6 (Mar 2018), 1299–1302. 1
- [WKR99] WEISKOPF D., KRAUS U., RUDER H.: Searchlight and doppler effects in the visualization of special relativity: a corrected derivation of the transformation of radiance. *ACM Trans. Graph.* 18, 3 (1999), 278–292. 3