A Framework for Transient Rendering

Adrian Jarabo¹

Julio Marco¹ Adolfo Muñoz¹ ¹Universidad de Zaragoza Raul Buisan¹ Wojciech Jarosz² ²Disney Research Zurich Diego Gutierrez¹



Figure 1: Our transient rendering framework allows time-resolved visualizations of light propagation. The caustic wavefront produced by the spherical lens above is distorted by the lens and also delayed due to the longer optical path traversed within the sphere. Note that in this scene we omit the propagation time of the last segment (from the scene to the camera).

Abstract

Recent advances in ultra-fast imaging have triggered many promising applications in graphics and vision, such as capturing transparent objects, estimating hidden geometry and materials, or visualizing light in motion. There is, however, very little work regarding the effective simulation and analysis of transient light transport, where the speed of light can no longer be considered infinite. We first introduce the transient path integral framework, formally describing light transport in transient state. We then analyze the difficulties arising when considering the light's time-of-flight in the simulation (rendering) of images and videos. We propose a novel density estimation technique that allows reusing sampled paths to reconstruct time-resolved radiance, and devise new sampling strategies that take into account the distribution of radiance along time in participating media. We then efficiently simulate time-resolved phenomena (such as caustic propagation, fluorescence or temporal chromatic dispersion), which can help design future ultra-fast imaging devices using an analysis-by-synthesis approach, as well as to achieve a better understanding of the nature of light transport.

CR Categories: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism—Raytracing

Keywords: transient rendering, transient light transport, importance sampling, bidirectional path tracing, progressive photon mapping

Links: 🗇 DL 🖾 PDF 🐻 WEB 🛇 VIDEO 📥 CODE

1 Introduction

One of the most general assumptions in computer graphics is to consider the speed of light to be infinite, leading to the simulation of light transport in steady state. This is a reasonable assumption, since most of the existing imaging hardware is very slow compared to the speed of light. Light transport in steady state has been extensively investigated in computer graphics (e.g. Dutré et al. [2006], Gutierrez et al. [2008a], Křivánek et al. [2013]), including for instance the gradient [Ramamoorthi et al. 2007; Jarosz et al. 2012] or frequency [Durand et al. 2005] domains. In contrast, work in the temporal domain has been mainly limited to simulating motion blur [Navarro et al. 2011] or time-of-flight imaging [Kolb et al. 2010].

We introduce in this paper a formal framework for transient rendering, where we lift the assumption of an infinite speed of light. While different works have looked into transient rendering [Smith et al. 2008; Jarabo 2012; Ament et al. 2014], they have approached the problem by proposing straight forward extensions of traditional steady-state algorithms, which are not adequate for efficient transient rendering for a variety of reasons. Firstly, the addition of the extra sampling domain given by the temporal dimension dramatically increases the convergence time of steady state rendering algorithms. Moreover, by extending the well-accepted path integral formulation [Veach 1997], we observe that paths contributing to each frame form a near-delta manifold in time, which makes sampling almost impossible. We solve this issue by devising new sampling strategies that improve the distribution of samples along the temporal domain, and a new density estimation technique that allows reconstructing the signal along time from such samples.

Our paper presents valuable insight apart from rendering applications. Recent advances in time-resolved imaging are starting to provide novel solutions to open problems, such as reconstructing hidden geometry [Velten et al. 2012a] or BRDFs [Naik et al. 2011], recovering depth of transparent objects [Kadambi et al. 2013], or even visualizing the propagation of light [Velten et al. 2013]. Despite these breakthroughs in technology, there is currently a lack of tools to efficiently simulate and analyze transient light transport. This would not only be beneficial for the graphics and vision communities, but it could open up a novel analysis-by-synthesis approach for applications in fields like optical imaging, material engineering or biomedicine as well. In addition, our framework can become

[©]ACM, 2014. This is the author's version of the work. It is posted here by permission of ACM for your personal use. Not for redistribution. The definitive version was published in ACM Transactions on Graphics, 33, 6, November 2014. doi.acm.org/10.1145/2661229.2661251

instrumental in teaching the complexities of light transport [Jarosz et al. 2012], as well as visualizing in detail some of its most cumbersome aspects, such as the formation of caustics, birefringence, or the temporal evolution of chromatic dispersion.

In particular, in this work we make the following contributions:

- Establishing a theoretical framework for rendering in transient state, based on the path integral formulation and including propagation in free space as well as scattering on both surfaces and in media. This allows us to analyze the main challenges in transient rendering.
- Developing a progressive kernel-based density estimation technique for path reuse that significantly improves the reconstruction of time-resolved radiance.
- Devising new sampling techniques for participating media to uniformly sample in the temporal domain, that complement traditional radiance-based sampling.
- Providing time-resolved simulations of several light transport phenomena which are impossible to see in steady state.

2 Related work

Transient radiative transfer. With advances in laser technology, capable of producing pulses of light in the order of a few femtoseconds, *transient* radiative transfer gained relevance in fields like optical imaging, material engineering or biomedicine. Many numerical strategies have been proposed, including Monte Carlo simulations, discrete ordinate methods, integral equation models or finite volume methods [Mitra and Kumar 1999; Zhang et al. 2013; Zhu and Liu 2013]. Often, these methods are applied on simplified scenarios with a particular application in mind, but a generalized framework has not yet been adopted.

Ultra-fast imaging. Several recent advances in ultra-fast imaging have found direct applications in computer graphics and vision. Raskar and Davis [2008] introduce the basic theoretical framework in light transport analysis that would later lead to a number of practical applications, such as reconstruction of hidden geometry [Kirmani et al. 2011; Velten et al. 2012a] or reflectance acquisition [Naik et al. 2011]. Velten et al. [2012b; 2013] have recently presented *femto-photography*, a technique that allows capturing time-resolved videos with an effective exposure time of one picosecond per frame, using a streak camera. Heide et al. [2013] later propose a cheaper setup using Photonic Mixing Devices (PMDs), while sacrificing temporal and spatial resolution. Kadambi and colleagues [2013] address multi path interference in time-of-flight sensors by recovering time profiles as a sequence of impulses, allowing them to recover depth from transparent objects.

Analysis of time-resolved light transport. Wu et al. [2012] analyze the propagation of light in the frequency domain, and show how the cross-dimensional transfer of information between the temporal and frequency domains can be applied to bare-sensor imaging. Later, Wu et al. [2013] used time-of-flight imaging to approximately decompose light transport into its different components of direct, indirect and subsurface illumination, by observing the temporal profiles at each pixel. Lin and colleagues [2014] perform a frequency-domain analysis of multifrequency time-of-flight cameras. Recently, O'Toole and colleagues [2014] derived transient light transport as a linear operator, as opposed to our formulation as a path integral, and showed how to combine the generation and acquisition of transient light transport for scene analysis. In this regard, our work can be seen as complementary: we provide a simulation (rendering) framework, suitable for an analysis-by-synthesis

approach to exploring novel ideas and applications, and to help better understand the mechanisms of light transport.

Transient rendering. The term *transient rendering* was first coined by Smith et al. [2008]. In their work, the authors generalize the rendering equation as a recursive operator including propagation of light at finite speed. The model provides a solid theoretical background for time-of-flight, computer vision applications, but does not provide a practical framework for transient rendering of global illumination. Keller et al. [2009] develop a time-of-flight sensor simulation, modeling the behavior of PMDs. These works are again geared towards time-of-flight applications; moreover, they are limited to surfaces, not taking into account the presence of participating media. Simulation of relativistic effects [Weiskopf et al. 1999; Jarabo et al. 2013] could also potentially benefit from our transient rendering framework.

Some recent works in computer graphics make use of transient state information: d'Eon and Irving [2011] quantize light propagation into a set of states, and model the transient state at each instant using Gaussians with variance proportional to time. These Gaussians are then integrated into the final image. The wave-based approach by Musbach et al. [2013] uses the Finite Difference Time Domain (FDTD) method to obtain a solution for Maxwell's equations, rendering complex effects like diffraction. In all these cases, however, the main goal is to render steady state images, not to analyze the propagation of light itself. Jarabo [2012] showed transient rendering results based on photon mapping and time-dependent density estimation, but limited to surfaces in the absence of participating media. Last, Ament et al. [2014] include time into the Radiative Transfer Equation in order to account for a continuously-varying indices of refraction in participating media, though they do not introduce efficient techniques for transient rendering.

Acoustic rendering. Our work is somewhat related to the field of acoustic rendering [Funkhouser et al. 2003]. Traditional light rendering techniques have been adapted to sound rendering, such as photon (*phonon*) mapping [Bertram et al. 2005] or precomputed *acoustic* radiance transfer [Antani et al. 2012]. Closest to our approach, the work by Siltanen et al. [2007] extends the radiosity method to include propagation delays due to the finite, though much slower, speed of sound. As opposed to us, they use finite element methods to compute sound transport, do not consider participating media, and do not propose sampling techniques for uniform temporal sample distribution.

3 Transient Path Integral Framework

We first extend the standard path integral formulation to transient state. This will allow us to formalize the notion of transient rendering, understand how to elevate steady state rendering to transient state, and, most importantly, identify the unique challenges of solving this more difficult light transport problem.

In the path integral formulation [Veach 1997], the image pixel intensity I is computed as an integral over the space of light transport paths Ω . For transient rendering, in addition to integrating over spatial coordinates, we must also integrate over the space of temporal delays ΔT of all paths:

$$I = \int_{\Omega} \int_{\Delta T} f(\overline{\mathbf{x}}, \overline{\Delta \mathbf{t}}) \,\mathrm{d}\mu(\overline{\Delta \mathbf{t}}) \,\mathrm{d}\mu(\overline{\mathbf{x}}), \tag{1}$$

where $\overline{\mathbf{x}} = \mathbf{x}_0 \dots \mathbf{x}_k$ represents the spatial coordinates of the k + 1 vertices of a length-k path with $k \ge 1$ segments. Vertex \mathbf{x}_0 lies on a light source, \mathbf{x}_k lies on the camera sensor, and $\mathbf{x}_1 \dots \mathbf{x}_{k-1}$ are



Figure 2: Spatio-temporal diagram of light propagation for a path with k = 2. Light is emitted at time t_0 , and reaches \mathbf{x}_1 at $t_0 + t(\mathbf{x}_0 \leftrightarrow \mathbf{x}_1)$. After a microscopic temporal delay Δt_1 , light emerges from \mathbf{x}_1 at t_1 and takes $t(\mathbf{x}_1 \leftrightarrow \mathbf{x}_2)$ time to reach \mathbf{x}_2 . The sensor may include a further temporal delay Δt_2 .

intermediate scattering vertices. The differential measure $d\mu(\overline{\mathbf{x}})$ denotes area integration for surfaces vertices and volume integration for media vertices. $\overline{\Delta \mathbf{t}} = \Delta t_0 \dots \Delta t_k$ defines a sequence of time delays and $d\mu(\overline{\Delta \mathbf{t}})$ denotes temporal integration at each path vertex.

We define the path contribution function $f(\overline{\mathbf{x}}, \overline{\Delta \mathbf{t}})$ as the original, but with the emission $L_{\rm e}$, path throughput \mathfrak{T} , and sensor importance $W_{\rm e}$ additionally depending on time:

$$f(\overline{\mathbf{x}}, \overline{\Delta \mathbf{t}}) = L_{\mathrm{e}}(\mathbf{x}_0 \to \mathbf{x}_1, \Delta t_0) \,\mathfrak{T}(\overline{\mathbf{x}}, \overline{\Delta \mathbf{t}}) \, W_{\mathrm{e}}(\mathbf{x}_{k-1} \to \mathbf{x}_k, \Delta t_k).$$
⁽²⁾

The temporal sensor importance W_e now defines not only the spatial and angular sensitivity, but also the region of time we are interested in evaluating. This could specify a delta function at a desired time, or more commonly, a finite interval of interest in the temporal domain (analogous to the shutter interval in steady state rendering, though at much smaller time scales). Likewise, the time parameter of the emission function L_e can define temporal variation in emission (e.g. pulses). The transient path throughput is now defined as:

$$\mathfrak{T}(\overline{\mathbf{x}}, \overline{\Delta \mathbf{t}}) = \left[\prod_{i=1}^{k-1} \rho(\mathbf{x}_i, \Delta t_i)\right] \left[\prod_{i=0}^{k-1} G(\mathbf{x}_i, \mathbf{x}_{i+1}) V(\mathbf{x}_i, \mathbf{x}_{i+1})\right].$$
(3)

Since we assume that the geometry is stationary (relative to the speed of light), the geometry and visibility terms depend only on the spatial coordinates of the path, as in steady state rendering. However, we extend the scattering kernel ρ with a temporal delay parameter Δt_i to account for potential time delays at each scattering vertex \mathbf{x}_i . Such delays can occur due to e.g. multiple internal reflections within micro-geometry [Westin et al. 1992], electromagnetic phase shifts in the Fresnel equations [Gondek et al. 1994; Sadeghi et al. 2012], or inelastic scattering effects such as fluorescence [Wilkie et al. 2001; Gutierrez et al. 2008b].

Time Delays. A transient light path is defined in terms of spatial and temporal coordinates. The temporal coordinates at each path vertex \mathbf{x}_i are t_i^- , the time immediately before the scattering event, and t_i , the time immediately after (see Figure 2). Both time coordinates can be obtained by accounting for all *propagation* delays between vertices $t(\mathbf{x}_i \leftrightarrow \mathbf{x}_{i+1})$ and *scattering* delays Δt_i at vertices along the path:

$$t_i^- = \sum_{j=0}^{i-1} \left(t(\mathbf{x}_j \leftrightarrow \mathbf{x}_{j+1}) + \Delta t_j \right), \quad t_i = t_i^- + \Delta t_i, \quad (4)$$

where t_0 and t_k denote the emission and detection times of a light path. The transient simulation is assumed to start at $t_0^- = 0$. In the



Figure 3: Left: The probability of finding a sample at a specific time instant (t_p or t_q) is nearly zero (Section 3). Middle: Density estimation on the temporal domain (Section 4) allows us to reconstruct radiance at any instant, although with varying bias and variance in time. **Right:** A more uniform distribution of samples in the temporal domain leads to more uniform bias and better reconstructions (Section 5).

general case of non-linear media [Gutierrez et al. 2005; Ihrke et al. 2007; Ament et al. 2014], propagation time along a path segment is:

$$t(\mathbf{x}_{j} \leftrightarrow \mathbf{x}_{j+1}) = \int_{s_{j}}^{s_{j+1}} \frac{\eta(\mathbf{x}_{r})}{c} \,\mathrm{d}r,\tag{5}$$

where *r* parametrizes the path of light between the two points, s_j and s_{j+1} are the parameters of the path at \mathbf{x}_j and \mathbf{x}_{j+1} , respectively, *c* is the speed of light in vacuum and $\eta(\mathbf{x}_r)$ represents the index of refraction of the medium at \mathbf{x}_r . In the typical scenario where η is constant along a path segment, Equation (5) reduces to a simple multiplication: $t(\mathbf{x}_j \leftrightarrow \mathbf{x}_{j+1}) = ||\mathbf{x}_j - \mathbf{x}_{j+1}|| \eta/c$. Figure 2 illustrates both the spatial and temporal dimensions of a path for the case of k = 2.

Numerical Integration. Similar to its steady state counterpart, the the transient path integral (1) can be numerically approximated using a Monte Carlo estimator:

$$\langle I \rangle = \frac{1}{n} \sum_{j=1}^{n} \frac{f(\overline{\mathbf{x}}_j, \overline{\Delta \mathbf{t}}_j)}{p(\overline{\mathbf{x}}_j, \overline{\Delta \mathbf{t}}_j)},\tag{6}$$

which averages n random paths $\overline{\mathbf{x}}_j$, $\overline{\Delta \mathbf{t}}_j$ drawn from a spatiotemporal probability distribution (pdf) $p(\overline{\mathbf{x}}_i, \overline{\Delta \mathbf{t}}_i)$ defined by the chosen path and time delay sampling strategy. In steady state, the pdf only needs to deal with the location of path vertices $\overline{\mathbf{x}}_i$.

3.1 Challenges of sampling in transient state

Equation (1) shows a new domain of scattering delays ΔT that must be sampled. Most existing path sampling techniques generate random paths incrementally, vertex-by-vertex, by locally importance sampling the scattering function ρ at each bounce, and optionally making deterministic shadow connections between light and camera subpaths. We could in principle elevate any such algorithm to transient state by simply sampling the transient scattering function $\rho(\mathbf{x}_i, \Delta t_i)$, instead of the steady state scattering function $\rho(\mathbf{x}_i)$.

Unfortunately, transient rendering poses hidden challenges, since *propagation* delays between vertices $t(\mathbf{x}_i \leftrightarrow \mathbf{x}_{i+1})$ are fundamentally different than *scattering* delays Δt_i defined at the light, sensor, and interior vertices. While scattering delays reside on a separate sampling domain ΔT , propagation delays are a direct consequence of the spatial positions of path vertices sampled from Ω . Hence, if spatial positions are determined by a steady state sampling routine ignorant of propagation delays, control of the propagation time in a path's total duration t_k is lost, leaving only the scattering delays Δt_i to control t_k .

Other factors resulting from the temporal structure of light transport make any naïve extension to transient rendering extremely inefficient: to visualize transient effects, the time window of both the sensor and the light source needs to be small ($\approx 10 \ picoseconds$); moreover, scattering events result in *femtosecond* temporal delays. The temporal domain of the path contribution thus becomes a near delta manifold (i.e. a *caustic* in time), which is virtually impossible to sample by random chance. Since the total path duration t_k cannot be directly controlled, deterministic shadow connections are rendered useless, having little chance of finding a non-zero contribution in both the light L_e and the sensor W_e . In general, the probability of randomly finding non-zero contribution for a specific time decreases as either Δt_i , L_e or W_e get closer to delta functions in the temporal domain, which are precisely the cases of interest in transient light transport.

When several distinct measurements of the path integral have to be computed, a common optimization strategy is to share randomly sampled paths to estimate all measurements simultaneously. This technique (path reuse) is utilized in the spatial domain in light tracing and bidirectional path tracing to estimate all pixels in the image plane at once. A similar situation occurs in the transient domain, where each frame f defines a specific sensor importance function $W_e^f(\mathbf{x}_{k-1} \rightarrow \mathbf{x}_k, t_k)$ and the time window covered by all frames is significantly larger than the per-frame time window. We could therefore leverage temporal path reuse to improve the efficiency of steady state path sampling methods when applied to rendering transient light transport. In practice, for every generated random path in Equation (6), we could evaluate the contribution functions for every frame f, which differ only in the temporal window of the sensor importance function W_e^f .

This path reuse technique is equivalent to histogram density estimation [Silverman 1986] in the *temporal domain* of the sensor, where each bin of the histogram represents one frame, and the bin's width h is the frame duration. Unfortunately, this type of density estimation produces very noisy results, especially for bins with very small width (i.e. exposure time). This results in a low convergence rate of $O(n^{-1/3})$ [Scott 1992], where n is the number of samples. This is illustrated in Figure 4: although obviously better than not reusing paths, results are still extremely noise even with a large amount of samples. Still, this suggests that more elaborated density estimation techniques may lead to better convergence rates and/or less noisy reconstructions.

In the following, we first show how kernel-based density estimation techniques in the temporal domain allow us to reconstruct radiance along time from a sparse set of samples (see Section 4 and Figure 3, middle). Then, we show how a skewed temporal sample distribution affects radiance reconstruction, and develop a set of sampling strategies for participating media that enable some control over propagation delays, leading to a more uniform distribution of samples in time and therefore more accuracy (see Section 5 and Figure 3, right).

4 Kernel-based temporal density estimation

Kernel-based density estimation is a widely known statistical tool to reconstruct a signal from randomly sampled values. These techniques significantly outperform histogram-based techniques (like the path reuse described above), especially for noisy data [Silverman 1986]. A kernel with finite bandwidth is used to obtain an estimate of the value of a signal at a given point by computing a weighted average of the set of random samples around such point. We thus introduce a temporal kernel K_T with bandwidth T to estimate incoming radiance I at the sensor at time t as a function of n samples of I:

$$\langle I_n \rangle = \frac{1}{n} \sum_{j=1}^n K_{\mathcal{T}}(\|t - t_{k,j}\|) \widehat{I}_j,$$
 (7)



Figure 4: Time-resolved irradiance computed at pixel (a) in the scene on the left using, no path reuse (green), histogram-based path reuse (red), and kernel-based path reuse (blue), for the same number of samples. Without path reuse it is extremely difficult to reconstruct the radiance, since the probability of finding a path arriving at the specific frame is close to zero. This is solved using path reuse, although with different levels of improvement: while histogram-based density estimation shows a very noisy result, our proposed progressive kernel-based estimation shows a solution with significantly lower variance, while preserving high-frequency features due to the progressive approach.

where $\widehat{I}_j = f(\overline{\mathbf{x}}_j, \overline{\Delta \mathbf{t}}_j)/p(\overline{\mathbf{x}}_j, \overline{\Delta \mathbf{t}}_j)$ is the contribution of path $\overline{\mathbf{x}}_j$ in the measured pixel, and $t_{k,j}$ is the total time of the path (4). Using this temporal density estimation kernel reduces variance, but at the cost of introducing bias (see Figure 3, middle). This can be solved by using consistent progressive approximations [Hachisuka et al. 2008; Knaus and Zwicker 2011], which converge to the correct solution in the limit.

Inspired by these works, we model our progressive density estimation along the temporal domain, for which we rely on the probabilistic approach for progressive photon mapping used by Knaus and Zwicker [2011]. We compute the estimate $\langle I_n \rangle$ in *n* steps, progressively reducing bias while allowing variance to increase; this is done by reducing the kernel bandwidth \mathcal{T} in each iteration as $\mathcal{T}_{j+1}/\mathcal{T}_j = (j+\alpha)/(j+1)$. The variance of our temporal progressive estimator vanishes with $O(n^{-\alpha})$ as expected, since the shrinking ratio is inversely proportional to the variance increase factor. Bias, on the other hand, vanishes with $O(n^{-2(1-\alpha)})$. Note that the parameter α defines the convergence of *both* sources of error (bias and variance). To find the optimal value that minimizes both, we use the *asymptotic mean square error* (AMSE), defined as:

$$AMSE(\langle I_n \rangle) = Var[I_n] + E[\epsilon_n]^2.$$
(8)

Using the convergence rate for both bias and variance, we find that the optimal α that minimizes the AMSE is $\alpha = 4/5$, which leads to a convergence of $O(n^{-4/5})$. This is significantly faster than using the histogram method, $O(n^{-1/3})$, which we illustrate in Figure 4. The detailed derivation of the behavior of the algorithm can be found in the supplementary material (Section B).

4.1 Transient progressive photon mapping

Our approach above is agnostic to the algorithm used to obtain the samples (e.g. samples in Figure 4 have been computed using path tracing). This means that it can be combined with biased density estimation-based algorithms such as (progressive) photon mapping [Jensen 2001; Hachisuka et al. 2008; Hachisuka and Jensen 2009], which is well suited for complex light paths such as spatial caustics. However, although using progressive photon mapping as the source of samples for our temporal density estimation is consistent in the limit, it results in suboptimal convergence due to the coupling of the bias and variance between the spatial and temporal kernels. Instead, we introduce the temporal domain into the photon mapping framework, by adding the temporal smoothing kernel $K_{\mathcal{T}}$ in the radiance estimation [Cammarano and Jensen 2002]. Radiance $\widehat{L}_{o}(\mathbf{x}, t)$ is estimated using M photons with contribution γ_i as

$$\widehat{L_{o}}(\mathbf{x},t) = \frac{1}{M} \sum_{i=1}^{M} K(\|\mathbf{x} - \mathbf{x}_{i}\|, \|t - t_{i}^{-}\|)\gamma_{i}.$$
 (9)

Combining both kernels into a single multivariate kernel allows controlling the variance increment in each step as a function of a single α , so that it increments at a rate of $(j + 1)/(j + \alpha)$, while reducing bias by progressively shrinking both the spatial and temporal kernel bandwidths (*R* and \mathcal{T} respectively). As shown in the supplementary material (Section C), these are reduced at each iteration *j* following:

$$\frac{\mathcal{T}_{j+1}}{\mathcal{T}_j} = \left(\frac{j+\alpha}{j+1}\right)^{\beta_{\mathcal{T}}}, \qquad \frac{R_{j+1}^2}{R_j^2} = \left(\frac{j+\alpha}{j+1}\right)^{\beta_R}, \qquad (10)$$

where β_T and β_R are scalars in the range [0, 1] controling how much each term is to be scaled separately, with $\beta_T + \beta_R = 1$. The convergence rate of the *combined* spatio-temporal density estimation is $O(n^{-4/7})^1$. Using this formulation allows us to handle complex light paths in transient state, while still progressively reducing bias and variance introduced by both progressive photon mapping and our temporal density estimation, in the spatial and temporal domains respectively. We refer to the supplementary material (Section C) for the detailed description of the algorithm, including the full derivation of the error and convergence rate.

5 Time Sampling in Participating Media

As we mentioned earlier (Section 3.1), the performance of our transient density estimation techniques can be further improved by a more uniform distribution of samples in time. This makes the relative error uniform in time and optimizes convergence (see Figure 3, right). Steady state sampling strategies aim to approximate radiance (path contribution). Since more radiant samples happen at earlier times (due to light attenuation), these sampling techniques skew the number of samples towards earlier times. As a consequence, there is a increase of error along time (see Figures 7 and 8). New sampling strategies are therefore needed for transient rendering.

Sampling strategies over scattering delays Δt_i have a negligible influence over the total path duration t_k (Figure 2). For surface rendering, scattering delays are the only control that sampling strategies can have on the temporal distribution of samples, and there is therefore little control over the total path duration. In participating media, however, sample points can be potentially located anywhere along the path of light, providing direct control also over the propagation times $t(\mathbf{x}_i \leftrightarrow \mathbf{x}_{i+1})$. In this section we develop new sampling strategies for participating media that target a uniform sample distribution in the time domain, by customizing:

- The pdf for each segment of the camera or light subpath (Section 5.1).
- The pdf for a shadow connection (connecting a vertex of the camera path to a vertex of the light path) via an additional vertex (Section 5.2).
- The pdf in the angular domain to obtain the direction towards the next interaction (Section 5.3).

Each of these sampling strategies ensures a uniform distribution of samples in time for each particular domain of the full path. Although this does not statistically ensure uniformity for the whole path, in practice the resulting distribution of total path duration t_k samples in time is close to uniform and therefore noise is reduced (the improvement over steady state strategies is discussed in Section 6). Note that these strategies are also agnostic of the properties of the media (except for the index of refraction), and can therefore be used in arbitrary participating media. Additionally, they can be combined with steady-state radiance sampling via multiple importance sampling (MIS) [Veach and Guibas 1995].

5.1 Sampling scattering distance in eye/light subpaths

Each of the segments of a subpath in participating media often shares the same steady-state sampling strategy, such as mean-freepath sampling, which does not necessarily ensure a uniform distribution of temporal location of vertices. We aim to find a pdf p(r)(where r is the scattering distance along one of the subpath segments) so that the probability distribution $p(\bigcup_{i=1}^{\infty}t_i)$ of temporal subpath vertex locations is uniform (see Figure 5, left). We first define $p(\bigcup_{i=1}^{\infty}t_i)$ based on the combined probability distribution $p(t_i)$ (temporal location of vertex \mathbf{x}_i in the light subpath) for all subpath vertices:

$$p\left(\bigcup_{i=1}^{\infty} t_i\right) = \sum_{i=1}^{\infty} p(t_i),\tag{11}$$

where $p(t_i)$ is recursively defined based on $p(t_{i-1})$. Given that $t_i = t(\mathbf{x}_i \leftrightarrow \mathbf{x}_{i-1}) + t_{i-1}$, as shown in Equation (4), we have

$$p(t_i) = \int_0^{t_i} p(t(\mathbf{x}_{i-1} \leftrightarrow \mathbf{x}_i)) p(t_{i-1}) dt_{i-1}, \qquad (12)$$

$$p(t_1) = p(t(\mathbf{x}_0 \leftrightarrow \mathbf{x}_1)), \tag{13}$$

since the probability of the addition of two random variables is the convolution of their probability distributions. $p(t(\mathbf{x}_i \leftrightarrow \mathbf{x}_{i-1}))$ is the probability distribution of the propagation time, which is related to the scattering distance pdf p(r) by a simple change of variable $r = \frac{c}{\eta} t(\mathbf{x}_{i-1} \leftrightarrow \mathbf{x}_i)$. Note that, in this notation, we are assuming (as previously discussed) that scattering delays Δt_i are negligible compared to propagation time. This definition is analogous for the eye subpath.

We show (see supplementary material, Section D.1) that the exponential distribution $p(r) = \lambda e^{-\lambda r}$ ensures that $p(\bigcup_{i=1}^{\infty} t_i)$ follows a uniform distribution for any λ parameter. Figure 6 (left) experimentally shows that this exponential distribution leads to this uniform probability for the whole subpath, while a uniform pdf leads to a non-uniform temporal sample distribution. In practice, λ modulates the average number of segments of the subpath: for a path ending at time t_e , the average number of segments with path duration $t_k \leq t_e$ is $\lambda_{\eta}^c t_e$. Our results show that an average of three or four vertices per subpath gives a good compromise between path length, efficiency and lack of correlation. Note that mean-free-path sampling is also an exponential distribution whose rate equals the extinction coefficient of the medium ($\lambda = \sigma_t$). Directly using mean-free-path sampling is thus optimal for time sampling when σ_t is close to the optimal λ .

Subpath termination. Russian roulette is a common strategy in steady state rendering algorithms. It probabilistically terminates subpaths at each scattering interaction, reducing longer paths with a small radiance contribution. In transient state, this unfortunately

¹Note that a naïve combination of the temporal (1D) and the spatial (2D) kernels would yield a slower convergence than the combined 3D kernel convergence $O(n^{-4/7})$ when using the optimal parameters $\alpha = 4/7$ and $\beta_T = 1/3$ reported in previous work [Kaplanyan and Dachsbacher 2013] (for volumetric density estimation) or in the statistics literature [Scott 1992].



Figure 5: Sampling strategies for participating media with a uniform distribution in the time domain. Left: Sampling scattering distance for a light subpath. This strategy can also be applied to eye subpaths. Middle Sampling shadow connections through a new indirect vertex: line-to-point sampling of shadow connections. Right: Sampling the angular scattering function (phase function).

translates into fewer samples as time advances, reducing the signalto-noise ratio (SNR) at higher frames. Instead, we simply terminate paths with a total duration greater than the established time frame.

While the temporal locations of subpath vertices are uniform, there is still little control over the spatial locations \mathbf{x}_i . These depend not only on scattering distances but also on scattering angles. As shadow rays are deterministic and depend on such spatial locations, uniformity cannot be ensured. To address this, we develop a new strategy that deals with such shadow connections (Section 5.2) and an angular sampling strategy (Section 5.3) that leads to an improved distribution in the temporal domain of the location-dependent propagation delays.

5.2 Sampling line-to-point shadow connections

Shadow rays are deterministic segments connecting a vertex in the eye subpath to another vertex in the light subpath, so their duration cannot be controlled. We introduce a new *indirect* shadow vertex whose position can be stochastically set to ensure a uniform sample distribution along the duration of the (extended) shadow connection. The geometry of this indirect connection is similar to equiangular sampling [Kalli and Cashwell 1977; Rief et al. 1984; Kulla and Fajardo 2012] (see Figure 5, middle).

Given a vertex \mathbf{x}_i of a light subpath, a vertex \mathbf{x}_{i+2} and a direction ω (importance sampled from the scattering function) on an eye subpath, our technique connects the two vertices via an indirect bounce at an importance-sampled location \mathbf{x}_{i+1} . If r_{i+1} and r_{i+2} are the distances from \mathbf{x}_{i+1} to \mathbf{x}_i and \mathbf{x}_{i+2} respectively, we importance sample r_{i+2} to enforce a uniform propagation time between the connected vertices { $\mathbf{x}_i, \mathbf{x}_{i+1}, \mathbf{x}_{i+2}$ }. This connection could also be done in reverse order (from \mathbf{x}_{i+2} to \mathbf{x}_i).

Given $\mathbf{l} = \mathbf{x}_i - \mathbf{x}_{i+2}$ and a connection time range (t_a, t_b) (in which we aim to get uniformly distributed samples), the pdf is:

$$p(r_{i+2}) = \frac{\eta}{c(t_b - t_a)} \left[1 + \frac{r_{i+2} - (\mathbf{l} \cdot \omega)}{\sqrt{r_{i+2}^2 - 2r_{i+2}(\mathbf{l} \cdot \omega) + (\mathbf{l} \cdot \mathbf{l})}} \right], (14)$$

which leads to the following inverse cumulative distribution function (cdf):

$$r_{i+2}(\xi) = \frac{(\xi(t_b - t_a) + t_a - t_i - \Delta t_{i+1})^2 - (\frac{\eta}{c})^2 (\mathbf{l} \cdot \mathbf{l})}{2\frac{\eta}{c} (\xi(t_b - t_a) + t_a - t_i - \Delta t_{i+1}) - 2(\frac{\eta}{c})^2 (\mathbf{l} \cdot \omega)}.$$
(15)

where $\xi \in [0, 1)$ is a random number. Assuming a rendered temporal range of $(0, t_e)$, we set the shadow connection limits to $t_a = t_i + t(\mathbf{x}_i \leftrightarrow \mathbf{x}_{i+2})$ and $t_b = t_e - \Delta t_k - \left(\sum_{j=i+2}^{k-1} t(\mathbf{x}_j \leftrightarrow \mathbf{x}_{j+1}) + \Delta t_j\right)$. The derivation of this pdf can be

found in the supplementary material (Section D.2). Figure 6 (middle) compares our line-to-point sampling strategy with other common strategies in terms of sample distribution along the temporal domain, leading to a uniform distribution of samples. Note that we discard all paths with a total duration longer than t_e (when $t_b < t_a$).

5.3 Angular sampling

Importance sampling the phase function generally leads again to a suboptimal distribution of samples in time. We propose a new angular pdf $p(\theta)$ to be applied at each interaction of the light subpath, which targets the temporal distribution of samples assuming that the next vertex \mathbf{x}_{i+1} casts a deterministic shadow ray towards the sensor. Given the sensor vertex \mathbf{x}_k and a sampled distance r_{i+1} between two consecutive vertices \mathbf{x}_i and \mathbf{x}_{i+1} (see Figure 5, right), this strategy ensures a uniform distribution of the total propagation time in { $\mathbf{x}_i, \mathbf{x}_{i+1}, \mathbf{x}_k$ }. The direction from \mathbf{x}_i to \mathbf{x}_{i+1} is $\omega = (\theta, \phi)$ (in spherical coordinates) where θ is the sampled angle and ϕ is uniformly sampled in [0..2 π). Note that the sampled angle θ is related to the direction towards the sensor ($\mathbf{l} = \mathbf{x}_k - \mathbf{x}_i$) and not to the incoming direction (which is often the system of reference for phase function importance sampling). This pdf is:

$$p(\theta) = \frac{r_{i+1}\sin\theta}{2\sqrt{r_{i+1}^2 + |\mathbf{l}|^2 - 2r_{i+1}|\mathbf{l}|\cos\theta}},$$
(16)

with the following inverse cdf:

$$\theta(\xi) = \arccos\left(\frac{|\mathbf{l}| - 2r_{i+1}^2\xi^2 - 2\xi r_{i+1}\left(|\mathbf{l}| - 1\right)}{r_{i+1}|\mathbf{l}|}\right).$$
 (17)

The supplementary material (Section D.3) contains the full derivation. This pdf prioritizes segments towards the target vertex \mathbf{x}_t , which helps in practice since backward directions often lead to paths that become too long for the rendered time frame. Figure 6 (right) shows how our angular sampling strategy leads to a uniform distribution of samples in time, as opposed to other alternatives. The shadow ray from vertex \mathbf{x}_{i+1} to the sensor in \mathbf{x}_k (and to every vertex in the eye subpath in bidirectional path tracing) is then cast by applying the sampling technique described in Section 5.2. Alternatively, the shadow ray could be cast from \mathbf{x}_i by applying MIS between this angular sampling and line-to-point time sampling (Section 5.2). We also apply the same angular sampling strategy for each interaction of the eye subpath, targeting the light source.

6 Results

Here we show and discuss our rendered scenes. For visualization we use selected frames of the animations; we refer the reader to the supplementary material for more rendered examples, and to the



Figure 6: Histogram of the number of samples along the temporal dimension for different sampling strategies. Left: Sample distribution for the whole light subpath, according to the importance sampling of subpath segments. Middle: Importance sampling of a line-to-point shadow connection Right: Angular importance sampling. Notice how our developed sampling strategies (exponential for segment sampling and the corresponding time sampling strategies in the other two cases) lead to a uniform distribution of samples along the temporal domain on each case. Both the line-to-point and the angular sampling are defined over a certain range.

video for the complete animations. In all the scenes light emission occurs at t = 0 with a delta pulse². Unless otherwise stated, we use transient path tracing and kernel-based density estimation (Section 4) for sampling and reconstruction, respectively. For the latter, we use a Perlin [2002] smoothing kernel, following previous work [Hachisuka et al. 2010; Kaplanyan and Dachsbacher 2013], with forty nearest neighbors to determine the initial kernel bandwidth. Unless noted otherwise, all results are shown in camera time [Velten et al. 2013] (i.e. including the propagation time of the last segment).

Figure 7 compares transient rendering using our three timebased sampling strategies (Section 5) against common radiancebased steady-state sampling techniques (mean-free-path and phasefunction sampling, and deterministic shadow connection). Our approach distributes samples more uniformly in time, which reduces variance along the whole animation, while significantly lowering noise in later frames. We obtain similar quality to standard sampling using two orders of magnitude less samples. These advantages are even more explicit when using our line-to-point sampling strategy to render single scattering, as shown in Figure 8, where we compare against equiangular sampling [Kulla and Fajardo 2012]. Figure 9 shows how the *combination* of our kernel-based density estimation and our time sampling strategies produces better results than using either technique in isolation.

Figures 1, S.3 and S.2 (the last two in the supplementary material) demonstrate the macroscopic delays due to traversing media with different orders of refraction, which leads to a temporal delay of the wavefront, especially visible in the caustics. In these examples, we use a transient version of the *photon beams* algorithm [Jarosz et al. 2011a] to obtain the radiance samples due to scattering in the media. For the single image visualization of Figure 1 we use the *peak-time* visualization proposed by Velten et al. [2013].



Figure 7: Comparison of our three time sampling strategies combined, against the standard techniques used in steady state, in the dragon scene accounting for multiple scattering (top). Each graph shows the time-resolved radiance (bottom) at pixel (a), for three different scattering coefficients $\sigma_s = \{0.2, 0.9, 1.5\}$, and absorption $\sigma_a = 0.1$. For 1K samples per pixel and frame, our combined techniques (red) feature a similar quality as standard steady state techniques with 128 times more samples (green), while with the same number of samples, our techniques significantly outperform standard sampling (blue), especially in highly scattering media. To emphasize the differences between sampling techniques, here we use the histogram path reuse (see Section 4). Additional results for other types of media can be found in the supplementary material.

Figure 10 compares our simulation against a real scene captured with the femtophotography technique of Velten et al. [2013]. We can see that our simulation faithfully reproduces the different orders of scattering events occurring during light propagation. Finally, Figure 11 shows different examples of non-trivial phenomena visible in transient state, including temporal chromatic dispersion due to wavelength-dependent index of refraction, refraction delays for ordinary and extraordinary rays in birefringent crystals [Weidlich and Wilkie 2008; Latorre et al. 2012] and fluorescence due to energy re-emission after absorption [Gutierrez et al. 2008b]. We refer to the supplementary video for full visualization of the different phenomena.

7 Discussion

In summary, we have extended the classical path space integral to include the temporal domain, and shown how the high frequency nature of transient light transport leads to severe sampling problems. We have proposed novel sampling strategies and density estimation techniques, which allow us to distribute samples uniformly in time, resulting in reduced variance and a constant distribution of bias. Our supplementary material contains a rigorous mathematical analysis of all our technical contributions. Last, we have presented simulations of interesting transient light transport effects using modified versions of a representative cross section of common rendering algorithms.

Apart from educational benefits, our work could be used to help design prototypes of novel ultra-fast imaging systems, or as a forward model for inverse problems such as recovering hidden geometry or material estimation. Our temporal progressive density estimation (Section 4) could also be used to accelerate radiance reconstruction in time-resolved imaging techniques, reducing the need for taking repeated measurements to improve the SNR. Moreover, synthetic ground truth data may become a very valuable tool for designing and benchmarking future ultra-fast imaging devices.

 $^{^{2}}$ We could use a Gaussian pulse, although this would introduce a number of downsides: 1) an ideal delta pulse does not introduce any additional temporal blur; 2) in reality, the scale of physical Gaussian pulses is 2-3 orders of magnitude smaller than the shutter open interval, in effect constituting a delta pulse; and 3) a delta pulse allows us to distinguish between effects caused by the actual behavior of light and effects due to limitations of current hardware.



Figure 8: Comparison of different sampling techniques for computing single scattering, in a scene consisting of a dragon illuminated by point light source within participating media (left). As opposed to simple mean-free-path sampling and the state-of-the-art equiangular sampling [Kulla and Fajardo 2012], that distributes samples based on radiance, our point-to-line sampling (Section 5.2) distributes samples so that the are uniformly distributed in time (bottom, right). This allows performing better in terms of relative error (bottom, left) when rendering time-resolved radiance, avoiding the radiance signal degradation at longer times. Here we use the histogram (Section 4) to emphasize the performance of the algorithms.

Our time-resolved simulations can help analyze the complex phenomena involved in light transport, and gain new insights. For instance, Figure 12 shows how during the early stages of light propagation, the first orders of scattering determine the shape of the light distribution (a spherical wavefront), but over time this shape becomes a Gaussian of increasing variance. This observation is consistent with previous work [Yoo and Alfano 1990], where it is shown that light in a medium exhibits diffusion after traveling about ten times the mean-free-path, and might explain some of the errors near the light source reported in the quantized diffusion model [D'Eon and Irving 2011]. This effect is more accentuated in the presence of anisotropic media, where the wavefront behavior is even more dominant.

Future work. There are many compelling avenues of future work: First, it would be interesting to extend a unified path sampling framework [Křivánek et al. 2014] to transient state. We have shown how the photon beams algorithm [Jarosz et al. 2011b] can be used in transient rendering, combined with our temporal density estimation; however, a spatio-temporal progressive photon beams framework would be needed to achieve optimal convergence in transient state. Additionally, by building a joint sampling strategy in both angle and distance, as in recent advanced steady state sampling techniques [Novák et al. 2012; Georgiev et al. 2013], we could leverage the benefits of both to ensure better uniformity in the temporal distribution of samples. Furthermore, the three proposed time sampling strategies are limited to participating media; extending this to surface transport results in a much narrower sampling space. Metropolis Light Transport techniques [Veach and Guibas 1997] represent promising candidates in this regard, where temporal mutation strategies would be needed.

We hope that our research will inspire future work on our understanding of light transport, the design of ultra-fast imaging and the development of novel rendering techniques. For instance, several geometric approaches to acoustic rendering are also based on ray tracing: a more extensive analysis of similarities between acous-



Figure 9: Selected frame of the dragon scene with $\sigma_s = 0.2$, rendered with a) standard sampling and histogram, b) our time sampling and histogram, c) standard sampling and our kernel-based density estimation, and d) time sampling and kernel-based density estimation. We can see how using our techniques combined lead to frames with significantly lower noise.



Figure 10: Comparison between the Cube scene from [Velten et al. 2013] and our rendered simulation of the same scene. Visible differences are due to approximate materials and camera properties.

tic and transient rendering might prove fruitful to both domains. Our code and datasets (scenes and movies) are publicly available at http://giga.cps.unizar.es/~ajarabo/pubs/transientSIGA14/.

Acknowledgments

We want to thank the reviewers for their insightful comments, and Mubbasir Kapadia for the link to sound rendering papers. This research has been partially funded by the European Commission, 7th Framework Programme, through projects GOLEM and VERVE, the Spanish Ministry of Economy and Competitiveness through project LIGHTSLICE, and project TAMA.

References

- AMENT, M., BERGMANN, C., AND WEISKOPF, D. 2014. Refractive radiative transfer equation. *ACM Trans. Graph.* 33, 2.
- ANTANI, L., CHANDAK, A., TAYLOR, M., AND MANOCHA, D. 2012. Direct-to-indirect acoustic radiance transfer. *IEEE Trans*actions on Visualization and Computer Graphics 18, 2.
- BERTRAM, M., DEINES, E., MOHRING, J., JEGOROVS, J., AND HAGEN, H. 2005. Phonon tracing for auralization and visualization of sound. In *IEEE Visualization '05*.
- CAMMARANO, M., AND JENSEN, H. W. 2002. Time dependent photon mapping. In *Eurographics Workshop on Rendering '02*.
- D'EON, E., AND IRVING, G. 2011. A quantized-diffusion model for rendering translucent materials. ACM Trans. Graph. 30, 4.



Figure 11: Examples of different phenomena observed in transient state: from left to right, temporal chromatic dispersion due to wavelength dependent index of refraction; ordinary and extraordinary image formation in a birefringent crystal; and energy re-emission in a fluorescent bunny. See the supplementary video for the full animations.



Figure 12: Time-resolved light transport from a point light source placed in the middle of an isotropic (left) and forward (right) scattering medium, emitting at time t = 0. Both media have a mean free path of 0.244 mm ($\sigma_t = 4.1 \text{ mm}^{-1}$). In the initial phase the light distribution is dominated by the wavefront shape of the low-order scattering events. In isotropic scattering, light distribution becomes Gaussian after traveling ten times the mean free path. In forward scattering, this distance is increased.

- DURAND, F., HOLZSCHUCH, N., SOLER, C., CHAN, E., AND SILLION, F. X. 2005. A frequency analysis of light transport. *ACM Trans. Graph.* 24, 3.
- DUTRÉ, P., BALA, K., AND BEKAERT, P. 2006. Advanced Global Illumination. AK Peters.
- FUNKHOUSER, T., TSINGOS, N., AND JOT, J.-M. 2003. Survey of methods for modeling sound propagation in interactive virtual environment systems. *Presence and Teleoperation*.
- GEORGIEV, I., KŘIVÁNEK, J., HACHISUKA, T., NOWROUZEZAHRAI, D., AND JAROSZ, W. 2013. Joint importance sampling of low-order volumetric scattering. ACM Trans. Graph. 32, 6.
- GONDEK, J. S., MEYER, G. W., AND NEWMAN, J. G. 1994. Wavelength dependent reflectance functions. In *SIGGRAPH '94*.
- GUTIERREZ, D., MUÑOZ, A., ANSON, O., AND SERON, F. 2005. Non-linear volume photon mapping. In *Eurographics Symposium on Rendering* '05.
- GUTIERREZ, D., NARASIMHAN, S. G., JENSEN, H. W., AND JAROSZ, W. 2008. Scattering. In ACM SIGGRAPH ASIA 2008 Courses.
- GUTIERREZ, D., SERON, F., MUÑOZ, A., AND ANSON, O. 2008. Visualizing underwater ocean optics. *Computer Graphics Forum* 27, 2.
- HACHISUKA, T., AND JENSEN, H. W. 2009. Stochastic progressive photon mapping. *ACM Trans. Graph.* 28, 5.
- HACHISUKA, T., OGAKI, S., AND JENSEN, H. W. 2008. Progressive photon mapping. *ACM Trans. Graph.* 27, 5.

- HACHISUKA, T., JAROSZ, W., AND JENSEN, H. W. 2010. A progressive error estimation framework for photon density estimation. ACM Trans. Graph. 29, 6.
- HEIDE, F., HULLIN, M., GREGSON, J., AND HEIDRICH, W. 2013. Low-budget transient imaging using photonic mixer devices. ACM Trans. Graph. 32, 4.
- IHRKE, I., ZIEGLER, G., TEVS, A., THEOBALT, C., MAGNOR, M., AND SEIDEL, H.-P. 2007. Eikonal rendering: Efficient light transport in refractive objects. ACM Trans. Graph. 26, 3.
- JARABO, A., MASIA, B., VELTEN, A., BARSI, C., RASKAR, R., AND GUTIERREZ, D. 2013. Rendering relativistic effects in transient imaging. In Congreso Español de Informática Gráfica (CEIG'13).
- JARABO, A. 2012. Femto-photography: Visualizing light in motion. Master's thesis, Universidad de Zaragoza.
- JAROSZ, W., NOWROUZEZAHRAI, D., SADEGHI, I., AND JENSEN, H. W. 2011. A comprehensive theory of volumetric radiance estimation using photon points and beams. ACM Trans. Graph. 30, 1.
- JAROSZ, W., NOWROUZEZAHRAI, D., THOMAS, R., SLOAN, P.-P., AND ZWICKER, M. 2011. Progressive photon beams. ACM Trans. Graph. 30, 6.
- JAROSZ, W., SCHÖNEFELD, V., KOBBELT, L., AND JENSEN, H. W. 2012. Theory, analysis and applications of 2D global illumination. ACM Trans. Graph. 31, 5.
- JENSEN, H. W. 2001. Realistic Image Synthesis Using Photon Mapping. AK Peters.
- KADAMBI, A., WHYTE, R., BHANDARI, A., STREETER, L., BARSI, C., DORRINGTON, A., AND RASKAR, R. 2013. Coded time of flight cameras: sparse deconvolution to address multipath interference and recover time profiles. ACM Trans. Graph. 32, 6.
- KALLI, H., AND CASHWELL, E. 1977. Evaluation of three Monte Carlo estimation schemes for flux at a point. Tech. Rep. LA-6865-MS, Los Alamos Scientific Lab, New Mexico, USA.
- KAPLANYAN, A. S., AND DACHSBACHER, C. 2013. Adaptive progressive photon mapping. ACM Trans. Graph. 32, 2.
- KELLER, M., AND KOLB, A. 2009. Real-time simulation of timeof-flight sensors. *Simulation Modelling Practice and Theory 17*, 5.
- KIRMANI, A., HUTCHISON, T., DAVIS, J., AND RASKAR, R. 2011. Looking around the corner using ultrafast transient imaging. *International Journal of Computer Vision 95*, 1.

- KNAUS, C., AND ZWICKER, M. 2011. Progressive photon mapping: A probabilistic approach. *ACM Trans. Graph.* 30, 3.
- KOLB, A., BARTH, E., KOCH, R., AND LARSEN, R. 2010. Timeof-flight sensors in computer graphics. *Computer Graphics Forum 29*, 1.
- KŘIVÁNEK, J., GEORGIEV, I., HACHISUKA, T., VÉVODA, P., ŠIK, M., NOWROUZEZAHRAI, D., AND JAROSZ, W. 2014. Unifying points, beams, and paths in volumetric light transport simulation. ACM Trans. Graph. 33, 4.
- KULLA, C., AND FAJARDO, M. 2012. Importance sampling techniques for path tracing in participating media. *Computer Graph ics Forum 31*, 4.
- KŘIVÁNEK, J., GEORGIEV, I., KAPLANYAN, A. S., AND CAÑADA, J. 2013. Recent advances in light transport simulation: theory & practice. In ACM SIGGRAPH 2013 Courses.
- LATORRE, P., SERON, F., AND GUTIERREZ, D. 2012. Birefringency: Calculation of refracted ray paths in biaxial crystals. *The Visual Computer* 28, 4.
- LIN, J., LIU, Y., HULLIN, M. B., AND DAI, Q. 2014. Fourier analysis on transient imaging with a multifrequency time-offlight camera. In *IEEE Conference on Computer Vision and Pattern Recognition '14*.
- MITRA, K., AND KUMAR, S. 1999. Development and comparison of models for light-pulse transport through scattering-absorbing media. *Appl Op 38*, 1.
- MUSBACH, A., MEYER, G. W., REITICH, F., AND OH, S. H. 2013. Full wave modelling of light propagation and reflection. *Computer Graphics Forum 32*, 6.
- NAIK, N., ZHAO, S., VELTEN, A., RASKAR, R., AND BALA, K. 2011. Single view reflectance capture using multiplexed scattering and time-of-flight imaging. ACM Trans. Graph. 30.
- NAVARRO, F., SERON, F., AND GUTIERREZ, D. 2011. Motion blur rendering: State of the art. *Computer Graphics Forum 30*, 1.
- NOVÁK, J., NOWROUZEZAHRAI, D., DACHSBACHER, C., AND JAROSZ, W. 2012. Virtual ray lights for rendering scenes with participating media. *ACM Trans. Graph.* 31, 4.
- O'TOOLE, M., HEIDE, F., XIAO, L., HULLIN, M. B., HEI-DRICH, W., AND KUTULAKOS, K. N. 2014. Temporal frequency probing for 5d transient analysis of global light transport. *ACM Trans. Graph. 33*, 4.
- PERLIN, K. 2002. Improving noise. ACM Trans. Graph. 21, 3.
- RAMAMOORTHI, R., MAHAJAN, D., AND BELHUMEUR, P. 2007. A first-order analysis of lighting, shading, and shadows. *ACM Trans. Graph.* 26, 1.
- RASKAR, R., AND DAVIS, J. 2008. 5d time-light transport matrix: What can we reason about scene properties? Tech. rep., MIT.
- RIEF, H., DUBI, A., AND ELPERIN, T. 1984. Track length estimation applied to point detector. *Nuclear Science and Engineering* 87.
- SADEGHI, I., MUÑOZ, A., LAVEN, P., JAROSZ, W., SERON, F., GUTIERREZ, D., AND JENSEN, H. W. 2012. Physically-based simulation of rainbows. ACM Trans. Graph. 31, 1.
- SCOTT, D. W. 1992. Multivariate Density Estimation: Theory, Practice, and Visualization. Wiley.

- SILTANEN, S., LOKKI, T., KIMINKI, S., AND SAVIOJA, L. 2007. The room acoustic rendering equation. *J. Acoust. Soc. Am. 122*, 3.
- SILVERMAN, B. W. 1986. Density Estimation for Statistics and Data Analysis. Taylor & Francis.
- SMITH, A., SKORUPSKI, J., AND DAVIS, J. 2008. Transient rendering. Tech. Rep. UCSC-SOE-08-26, School of Engineering, University of California, Santa Cruz.
- VEACH, E., AND GUIBAS, L. J. 1995. Optimally combining sampling techniques for Monte Carlo rendering. In SIGGRAPH '95.
- VEACH, E., AND GUIBAS, L. J. 1997. Metropolis light transport. In *SIGGRAPH* '97.
- VEACH, E. 1997. *Robust Monte Carlo methods for light transport simulation*. PhD thesis, Stanford.
- VELTEN, A., WILLWACHER, T., GUPTA, O., VEERARAGHAVAN, A., BAWENDI, M. G., AND RASKAR, R. 2012. Recovering three-dimensional shape around a corner using ultrafast time-offlight imaging. *Nature Communications*, 3.
- VELTEN, A., WU, D., JARABO, A., MASIA, B., BARSI, C., LAWSON, E., JOSHI, C., GUTIERREZ, D., BAWENDI, M. G., AND RASKAR, R. 2012. Relativistic ultrafast rendering using time-of-flight imaging. In ACM SIGGRAPH 2012 Talks.
- VELTEN, A., WU, D., JARABO, A., MASIA, B., BARSI, C., JOSHI, C., LAWSON, E., BAWENDI, M., GUTIERREZ, D., AND RASKAR, R. 2013. Femto-photography: Capturing and visualizing the propagation of light. ACM Trans. Graph. 32, 4.
- WEIDLICH, A., AND WILKIE, A. 2008. Realistic rendering of birefringency in uniaxial crystals. ACM Trans. Graph. 27, 1.
- WEISKOPF, D., KRAUS, U., AND RUDER, H. 1999. Searchlight and doppler effects in the visualization of special relativity: a corrected derivation of the transformation of radiance. *ACM Trans. Graph. 18*, 3.
- WESTIN, S. H., ARVO, J. R., AND TORRANCE, K. E. 1992. Predicting reflectance functions from complex surfaces. In *SIG*-*GRAPH* '92.
- WILKIE, A., TOBLER, R. F., AND PURGATHOFER, W. 2001. Combined rendering of polarization and fluorescence effects. In *Eurographics Workshop on Rendering Techniques '01.*
- WU, D., WETZSTEIN, G., BARSI, C., WILLWACHER, T., O'TOOLE, M., NAIK, N., DAI, Q., KUTULAKOS, K., AND RASKAR, R. 2012. Frequency analysis of transient light transport with applications in bare sensor imaging. In *European Conference on Computer Vision '12*.
- WU, D., VELTEN, A., O'TOOLE, M., MASIA, B., AGRAWAL, A., DAI, Q., AND RASKAR, R. 2013. Decomposing global light transport using time of flight imaging. *International Journal of Computer Vision 105*, 3.
- YOO, K. M., AND ALFANO, R. R. 1990. Time-resolved coherent and incoherent components of forward light scattering in random media. *Opt. Lett.* 15, 6.
- ZHANG, Y., YI, H., AND TAN, H. 2013. One-dimensional transient radiative transfer by lattice Boltzmann method. *Optics Express 21*, 21.
- ZHU, C., AND LIU, Q. 2013. Review of Monte Carlo modeling of light transport in tissues. J. Biomed. Opt. 18, 5.