



DARTS: Diffusion Approximated Residual Time Sampling for Time-of-flight Rendering in Homogeneous Scattering Media

QIANYUE HE, DONGYU DU, and HAITIAN JIANG, Tsinghua Shenzhen International Graduate School, China
XIN JIN*, Tsinghua Shenzhen International Graduate School, China

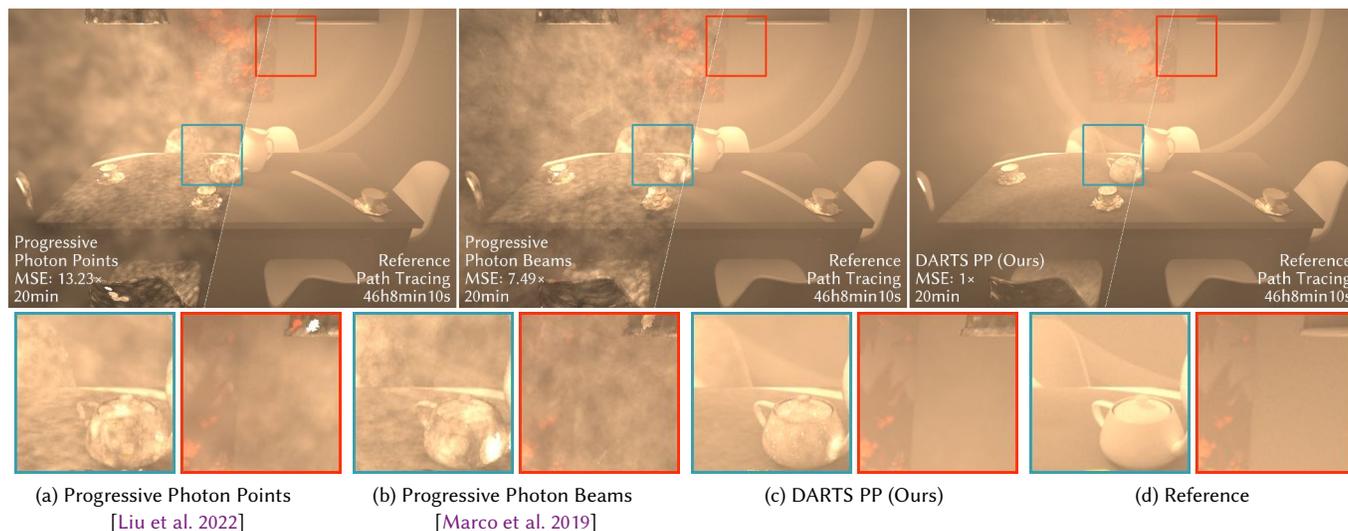


Fig. 1. Rendering of time-resolved transport using proposed DARTS in a scene with complex surface materials and homogeneous scattering media. DARTS integrates transient diffusion approximation into the path construction and adapts our elliptical sampling to provide path length control, enabling high quality time-of-flight rendering and can be compatible with different existing frameworks. The example scene is illuminated by two non-synchronized pulse emitters with different start times of emission. Each image is rendered for 20 minutes. It can be seen that our sampling approach can greatly improve the SOTA photon based methods and provide lower overall MSE in the same rendering time.

Time-of-flight (ToF) devices have greatly propelled the advancement of various multi-modal perception applications. However, achieving accurate rendering of time-resolved information remains a challenge, particularly in scenes involving complex geometries, diverse materials and participating media. Existing ToF rendering works have demonstrated notable results, yet they struggle with scenes involving scattering media and camera-warped settings. Other steady-state volumetric rendering methods exhibit significant bias or variance when directly applied to ToF rendering tasks. To address these challenges, we integrate transient diffusion theory into path construction and propose novel sampling methods for free-path distance and scattering direction, via resampled importance sampling and offline

tabulation. An elliptical sampling method is further adapted to provide controllable vertex connection satisfying any required photon traversal time. In contrast to the existing temporal uniform sampling strategy, our method is the first to consider the contribution of transient radiance to importance-sample the full path, and thus enables improved temporal path construction under multiple scattering settings. The proposed method can be integrated into both path tracing and photon-based frameworks, delivering significant improvements in quality and efficiency with at least a 5x MSE reduction versus SOTA methods in equal rendering time.

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

Additional Key Words and Phrases: transient rendering, time-gated cameras, participating media, modeling and simulation

ACM Reference Format:

Qianyu He, Dongyu Du, Haitian Jiang, and Xin Jin. 2024. DARTS: Diffusion Approximated Residual Time Sampling for Time-of-flight Rendering in Homogeneous Scattering Media. *ACM Trans. Graph.* 43, 6, Article 240 (December 2024), 14 pages. <https://doi.org/10.1145/3687930>

1 Introduction

The past decade has witnessed noteworthy advancements in time-of-flight (ToF) imaging methods which have the capability to capture transient responses of propagating photons utilizing ultra-fast sensors [Jarabo et al. 2017]. With the incorporation of temporal information, ToF imaging systems have revolutionized conventional

*Corresponding Author

Authors' Contact Information: Qianyu He, he-qy22@mails.tsinghua.edu.cn; Dongyu Du, dudy19@mails.tsinghua.edu.cn; Haitian Jiang, jht22@tsinghua.org.cn, Tsinghua Shenzhen International Graduate School, China; Xin Jin, jin.xin@sz.tsinghua.edu.cn, Tsinghua Shenzhen International Graduate School, China.

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

© 2024 Copyright held by the owner/author(s). Publication rights licensed to ACM. ACM 1557-7368/2024/12-ART240 <https://doi.org/10.1145/3687930>

imaging, making it possible to sense targets beyond the line of sight [Royo et al. 2023] and operate well in extremely challenging environment, such as heavy fog [Du et al. 2022]. Consequently, ToF devices propel the practical process of autonomous driving, robotic perception and scientific exploration, etc.

In tandem with ToF imaging techniques, ToF rendering provides an enhanced understanding of temporal data across various scene configurations, and thus contributes to the development of optical transmission theories and imaging algorithms, and the optimization of sensor systems [Marco et al. 2017a; Zhang et al. 2022]. Moreover, ToF rendering plays a crucial role in generating extensive ToF datasets across diverse scenarios involving complex participating media, geometries and material properties, which are imperative for the advancement of data-driven methods [Attal et al. 2021].

However, challenges reside in achieving efficient ToF rendering within participating media due to the complexity of temporal path construction. Despite substantial efforts directed towards steady-state scattering rendering [Herholz et al. 2019; Lin et al. 2021], the spatial sampling strategies fall short in generating temporal scattering paths. For transient rendering methods, certain methods achieve temporal uniformity in path samples [Jarabo et al. 2014] or populate the sampling domain more densely [Marco et al. 2019], but their accuracy is limited by the neglect of importance-sampling the contribution of transient radiance. Achieving effective importance sampling requires that major volumetric sampling processes, such as distance and direction sampling, incorporate transient radiance information. This requirement poses a challenge as the transient radiance information depends on global adjoint transport rather than just local transport functions. Additionally, the challenge persists in effectively imposing path time (length) constraints. While the line-to-point sampling method devised by Jarabo et al. [2014] can impose ellipsoidal path length constraint in participating media, it prioritizes the uniform distribution of path samples over effectively controlling path lengths. Other existing methods to impose path length constraints either prove not directly applicable for participating media [Pediredla et al. 2019], or inefficient for full transport and camera-warped settings [Liu et al. 2022], where scene-to-sensor transport time must be considered [Velten et al. 2013].

In this paper, we propose a diffusion approximated residual time sampling method (DARTS, for short), which provides full transient path construction and effective vertex connection within complex volumetric scenes and under camera-warped settings. To address the challenge of constructing effective temporal sampling paths in scattering media, instead of adopting uniform sampling in the time domain [Jarabo et al. 2014; Marco et al. 2019], which overlooks the radiance contribution differences of different path samples, we perform importance sampling on the transient radiance by integrating the transient diffusion approximation (DA) into the rendering process. This approach allows us to obtain improved free-path distance and direction samples with enhanced overall radiance, leading to better convergence performance. To impose path time constraints, we extend the ellipsoidal connection method and further combine it with the proposed DA-based direction sampling to bypass challenges introduced by the reparameterization of ellipsoidal connection [Pediredla et al. 2019] and avoid sampling

inefficiency in camera-warped settings [Liu et al. 2022]. The proposed method is inherently unbiased and introduces negligible extra memory overhead compared to the naive methods.

In particular, we make the following contributions:

- We *propose* a novel distance sampling method named DA distance sampling based on transient diffusion theory.
- We *propose* a novel direction sampling method named elliptical DA direction sampling, by tabulating the transient DA values integrated in an equal-time ellipse and *develop* corresponding multiple importance sampling (MIS) strategy.
- We *extend* the ellipsoidal connection to volumetric rendering and enable its capability for importance sampling and direction reuse, to effectively control the path length.
- We *demonstrate* the proposed method outperforms the existing method with at least a 5x MSE reduction and can be integrated as straightforward plug-ins for both path tracing and photon based frameworks.

The code for the proposed method in both *pbirt-v3*¹ [Pharr et al. 2023] and *Tungsten*² [Bitterli 2018] frameworks is provided in our supplementary materials.

2 Related Work

2.1 Time-of-flight imaging devices

ToF devices employ ultra-fast sensors to capture and count the photons received at different time points [Jarabo et al. 2017]. This time-resolved information is recorded to generate single time-gated image or sequences of temporal waveforms with high temporal resolution [Jarabo 2012], which can be applied in imaging through scattering media [Du et al. 2022; Wu et al. 2018] and obscurants [Halimi et al. 2021; Kijima et al. 2021], non-line-of-sight imaging [Faccio et al. 2020; Rapp et al. 2020; Xin et al. 2019], material estimation [Shem-Tov et al. 2020; Zickus et al. 2020] and improved depth sensing [Gruber et al. 2019; Walia et al. 2022].

2.2 Time-of-flight rendering

Transient rendering. It aims to simulate the temporal responses as a sequence of frames. Existing works can guarantee physical correctness, from the basics of transient light transport for both forward [Jarabo et al. 2014] and differentiable [Yi et al. 2021] cases, to more complex vector light transport [Jarabo and Arellano 2018] and in media with spatially various refraction index [Ament et al. 2014], or to directly sample the three-bounce approximation in the non-line-of-sight scenario [Iseringhausen and Hullin 2020]. However, the convergence for transient rendering is slow due to the lack of efficient sampling method [Bitterli 2016b]. Efforts have been made to importance-sample the uniform distribution of path lengths for better temporal density estimation and path reuse [Jarabo et al. 2014]. Subsequent methods [Marco et al. 2019, 2017b] extend the photon beams method to transient state, but they introduce bias in exchange for reduced variance. The above methods heavily rely on temporal path reuse, making it challenging to be extended to time-gated rendering where temporal path reuse will lead to massive sample

¹<https://github.com/mmp/pbirt-v3>

²<https://github.com/GhostatSpirit/tungsten-transient-public>

rejection. Also, existing works often disregard the importance sampling of transient radiance contributions during *path construction* and therefore cannot guarantee a good approximation of the integrand for Monte Carlo integration. Other works propose different insights like using instant radiosity [Pan et al. 2019] for fast rendering, while the visual effects are constrained to diffuse scattering, or employ sampling methods biased towards receivers [Lima et al. 2011; Periyasamy and Pramanik 2016] and consider primarily the spatial distribution of radiance, therefore neglecting the temporal distribution.

Time-gated rendering. Researchers focus on improving path connections to impose constraints on the path time (length). Pediredla et al. [2019] demonstrate that, given two vertices to be connected and a target path length, an ellipsoid can be defined and any connection with one intermediate vertex sampled on this ellipsoid has equal path length. Unfortunately, their parameterization cannot be directly extended to volumetric rendering, as volumetric sampling lacks intersecting polygons. Additionally, parameterizing the sampling space around the ellipsoid center prevents effective importance sampling and the reuse of sampled directions. Other methods have taken steps to populate the sampling space more densely. Liu et al. [2022] extend the methods introduced by Deng et al. [2019] to the time domain, illustrating that time-related sampling can be viewed as sampling from the sliced high-dimensional photon primitives, where dense samples can be directly drawn from. However, this method is not efficient under camera-warped settings [Velten et al. 2013] or with surface transport. These methods in general focus on improving *path connection* for better path length control and do not yield better multiple scattering paths, which are crucial for high-order scattering scenarios and temporal importance sampling.

2.3 Rendering in homogeneous scattering media

Various approaches have been developed [Novák et al. 2018] to improve the rendering quality in homogeneous scattering media. We will focus on two most widely adopted types and briefly outline the challenges of extending these methods to the time domain.

2.3.1 Density Estimation Based Methods. We refer to these methods as *photon based methods*. Stemming from photon mapping [Jensen and Christensen 1998], this kind of two-pass biased estimator is later upgraded to progressive ones [Hachisuka and Jensen 2009]. Jarosz et al. [2011] extend the point-point 3D kernel estimator to various point-beam and beam-beam estimators, which are later combined with bidirectional methods [Křivánek et al. 2014]. Subsequent works [Bitterli and Jarosz 2017; Deng et al. 2019] devise unbiased photon estimators by shrinking the kernel to a spatial delta-function to eliminate bias. Though these methods has temporal extensions [Liu et al. 2022; Marco et al. 2017b], they currently prove to be either ineffective for reduced rendering time range, such as time-gated rendering, or inefficient for camera-warped settings. This inefficiency arises from the inability to importance sample transient radiance, as well as the dependence on temporal path reuse [Jarabo et al. 2014].

2.3.2 Monte Carlo Path Tracing. Previous methods aim to analytically approximate terms in the integrand during distance or direction sampling [Georgiev et al. 2013; Kulla and Fajardo 2012]. Recent methods based on path guiding [Herholz et al. 2016; Vorba

et al. 2014] employ online learning to fit the local radiance distribution. Herholz et al. [2019] further combine path guiding with zero variance random walk theory [Hoogenboom 2008; Ren et al. 2008] to guide all scattering events. However, these methods are inherently time-agnostic, and due to the curse of dimensionality, extending online-learning based methods directly to the time domain significantly increases training samples sparsity.

Moreover, neither of the above methods adequately addresses the importance of more effective *path construction*, which is crucial for rendering high-order scattering. Therefore, we aim to establish a unified framework capable of both time-gated and transient rendering through optimizing both path construction and path connection strategies.

3 Background

Time-of-flight renderers lift the infinite speed of light assumption. Thus, the radiance transport theory and estimators should explicitly account for path time information. In the transient path integral framework proposed by Jarabo et al. [2014], the intensity I of an image pixel is given by:

$$I = \int_{\Delta t_0} \int_{\Omega} L_e(\mathbf{x}_e \rightarrow \mathbf{x}_{k-1}, t_e, \Delta t_e) f(\bar{\mathbf{x}}) W(\mathbf{x}_1 \rightarrow \mathbf{x}_0, \|\bar{\mathbf{x}}\|) d\mu(\bar{\mathbf{x}}) d\Delta t_e, \quad (1)$$

where $\bar{\mathbf{x}} = \mathbf{x}_0 \mathbf{x}_1 \dots \mathbf{x}_{k-1} \mathbf{x}_e$ is the simulated path with $k + 1$ vertices; $\mathbf{x}_i \in \mathbb{R}^3$ denotes the position of path vertices; \mathbf{x}_0 and \mathbf{x}_e denote sensor and emitter vertex, respectively; $d\mu(\cdot)$ is the Lebesgue measure; Ω denotes path space. $W(\mathbf{x}_1 \rightarrow \mathbf{x}_0, \|\bar{\mathbf{x}}\|)$ is the response function of the sensor and is most relevant to temporal transport since it is the function of path optical length $\|\bar{\mathbf{x}}\|$. Since our work does not account for fluorescence and other microscopic scattering delays, factors that affect $\|\bar{\mathbf{x}}\|$ can only be the accumulated path length and emission duration Δt_e . $L_e(\mathbf{x}_k \rightarrow \mathbf{x}_{k-1}, t_e, \Delta t_e)$ is the emission term where t_e denotes the start time of emission and Δt_0 denotes the integral space of Δt_e . Typically, for a simple pulse emitter with the start time of emission t_e set to 0, the simplified term L_e can encompass most of the use cases. In the following, we will ignore the view dependence in $\mathbf{x}_1 \rightarrow \mathbf{x}_0$ for simplicity. $f(\bar{\mathbf{x}})$ is path throughput term and takes the following form:

$$f(\bar{\mathbf{x}}) = g(\bar{\mathbf{x}}) G(\bar{\mathbf{x}}), \quad (2)$$

$G(\bar{\mathbf{x}})$ is the throughput term consisting of path transmittance, measure conversion terms and visibility terms; $g(\bar{\mathbf{x}})$ denotes the local throughput function, which is the product of several bidirectional scattering distribution function (BSDF) terms and phase function terms. In scenes filled with homogeneous participating media with constant relative refraction index η , the transport time for a path $\bar{\mathbf{x}}$ with k vertices can be simplified to Equation (3):

$$t(\bar{\mathbf{x}}) = \frac{\eta}{c} \|\bar{\mathbf{x}}\| = \frac{\eta}{c} \sum_{i=0}^{k-1} \|\mathbf{x}_i - \mathbf{x}_{i+1}\|, \quad (3)$$

c represents speed of light in vacuum. Note that path length $\|\bar{\mathbf{x}}\|$ has the same meaning as path time $t(\bar{\mathbf{x}})$ and only differs by a scalar scaling factor. Therefore, path length $\|\bar{\mathbf{x}}\|$ and path time $t(\bar{\mathbf{x}})$ will be used interchangeably, with a slight abuse of notation.

Equation (1) can be estimated numerically using Monte Carlo integration. A single path with multiple vertices originating from camera \mathbf{x}_0 is traced and each vertex is connected to the emitter, as shown in Figure 2. The traced path, direct shadow connection and generalized shadow connection (connection with extra control vertices) are represented by solid black line, dashed yellow line and dashed red lines, respectively. The transient radiance estimator for the estimation of the integral defined by Equation (1) takes the form of the following [Pediredla et al. 2019]:

$$\tilde{I} = \frac{1}{N} \sum_{n=1}^N \frac{W(\|\overline{\mathbf{x}}_n\|)f(\overline{\mathbf{x}}_n)L_e}{p(\overline{\mathbf{x}}_n)}, \quad (4)$$

where N different paths $\overline{\mathbf{x}}_n, n = 1, 2, \dots, N$ are used to estimate the pixel intensity; $p(\overline{\mathbf{x}}_n)$ denotes the sampling probability density function (PDF) to the joint distribution of $\overline{\mathbf{x}}_n$. It can be seen that the temporal response function $W(\|\overline{\mathbf{x}}_n\|)$ is the only different part from the steady-state estimator.

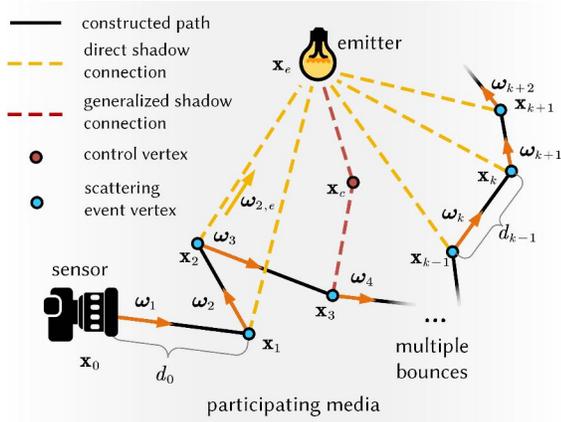


Fig. 2. Unidirectional path tracing in a scene filled with participating media. Random walk generates a path (solid black lines) with multiple vertices and all the vertices are connected to the emitter either through direct shadow connection (dashed yellow lines) or generalized shadow connection with control vertices (dashed red lines)

The high variance in transient radiance estimators arises primarily from two factors: the omission of radiance contributions in importance sampling, leading to less effective multiple scattering paths, and the inefficiency in imposing path length constraints. Detailed variance analysis can be found in our supplementary note (Section A.1). In Section 4, we address the first issue by using the diffusion approximation (DA) to enhance importance sampling for distance sampling. In Section 5, we tackle the second issue by combining the ellipsoidal connection methodology with DA, improving direction sampling and effectively imposing path length constraints.

4 Diffusion Approximated Distance Sampling

4.1 Residual Time Estimator

Different from existing methods that exhibit limited awareness of transient radiance distribution, our method aims to construct paths where all vertices are importance-sampled according to radiance

contribution of the specific target time intervals. Denoting target time for the full path as T_t and the time taken for photons travelling to the k -th vertex as elapsed time $T_{e,k}$, the *residual time* $T_{res,k}$ for the k -th vertex can be defined as:

$$T_{res,k} = T_t - T_{e,k} = T_t - \frac{\eta}{c} \sum_{i=0}^{k-1} \|\mathbf{x}_i - \mathbf{x}_{i+1}\|, \quad (5)$$

We note that residual time T_{res} is usually longer than the path time for direct connection, and imposing path length constraints cannot be effectively achieved through naive direct connection.

Due to the recursive nature of the path construction process, the optimal sampling procedure at each vertex involves first importance sampling the ray direction based on the transient radiance and then the residual time, aiming to maximize the average radiance contribution along the sampled distance. Once the direction is determined, a free-path distance is sampled along this direction to establish a new vertex for path connection. Since direction sampling applies to both path construction and path connection, it will be discussed in detail in Section 5.1. This section focuses on distance sampling given a specific ray direction.

To simplify the mathematical representation in Equation (4), we formulate the transient radiance estimator in a recursive form. We decompose the incident radiance of vertex \mathbf{x}_{k+1} into direct and indirect components, as depicted in Figure 3. The combination of direct and indirect components at vertex \mathbf{x}_{k+1} is the indirect radiance incident at \mathbf{x}_k before exhibiting volumetric attenuation. The incident indirect radiance can then be defined with respect to the residual time:

$$\tilde{L}_k(\mathbf{x}_k, -\omega_{k+1}, T_{res,k}) = \underbrace{\left(W(\|\overline{\mathbf{x}}'_{k+1}\|)L_d(\mathbf{x}_{k+1}, T_{res,k+1}) \right)}_I + \underbrace{\left(\frac{f_m(\omega_k, \omega_{k+1})}{p_d(\omega_{k+1})} \tilde{L}_{k+1} \right)}_{II} \frac{\sigma_s \exp(-\sigma_e d_k)}{p_t(d_k)}, T_{res,k} = T_{res,k+1} + \frac{\eta}{c} d_k, \quad (6)$$

$\overline{\mathbf{x}}'_{k+1}$ denotes a path where the complete sensor-to-emitter path is formed by connecting vertex \mathbf{x}_{k+1} to \mathbf{x}_e , possibly via intermediate vertices. σ_s, σ_e denotes scattering and extinction coefficients, respectively. $-\omega_{k+1}$ denotes the exiting direction at \mathbf{x}_{k+1} , and the sampling method of this direction will be discussed in Section 5.1. Part (I) in Equation (6) is the exiting direct radiance at vertex \mathbf{x}_{k+1} , and here we use word *direct* to denote connection paths produced by sampling control vertices before the sampled emitter, with a slight abuse of notation. Given the path length target T_t , valid sample must have $T_{res,k+1}$ as ToF to account for elapsed time of \mathbf{x}_k and transport time for distance d_k . Likewise, for the indirect incident radiance \tilde{L}_{k+1} in part (II), radiance sample with subsequent scattering should be able to satisfy the path length constraint, therefore it should have $T_{res,k+1}$ as ToF as well. As \tilde{L}_{k+1} is sampled in scattering direction, the Monte Carlo term is applied to convert it to estimated exiting radiance. All components will be attenuated due to scattering and absorption within the transport distance d_k produced by free-path distance sampling.

Since both direct and indirect component have the same time $T_{\text{res},k+1}$, Equation (6) can be compressed to a unified residual time radiance representation that encompasses both direct and indirect components as L' , as shown in Figure 3:

$$\tilde{L}_k(\mathbf{x}_k, -\boldsymbol{\omega}_{k+1}, T_{\text{res},k}) = \frac{\sigma_s \exp(-\sigma_e d_k) L'(\mathbf{x}_{k+1}, -\boldsymbol{\omega}_{k+1}, T_{\text{res},k+1})}{p_t(d_k)}, \quad (7)$$

To achieve low variance estimation, we need to find a sampling distribution whose PDF takes the following form:

$$p_t(d_k) = \frac{\sigma_s \exp(-\sigma_e d_k) L'(\mathbf{x}_{k+1}, -\boldsymbol{\omega}_{k+1}, T_{\text{res},k+1})}{Z}, \quad (8)$$

where Z is the normalizing constant. The key point then comes down to sample the full paths according to the attenuated transient radiance $\exp(-\sigma_e d_k) L'(\mathbf{x}_{k+1}, -\boldsymbol{\omega}_{k+1}, T_{\text{res},k+1})$, in order to calculate the next vertex position \mathbf{x}_{k+1} . Equation (8) underscores the intuition behind using the unified residual time in Equation (7): sampling should be concentrated in areas where the combined direct and indirect components are substantial.

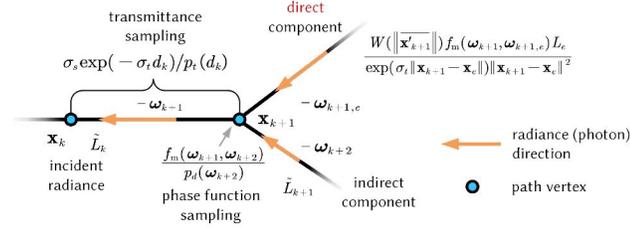


Fig. 3. Recursive decomposition of indirect radiance \tilde{L}_k into the direct and indirect components. Through this decomposition, the estimator can actually be written as the summation of an infinite series of direct component at each vertex. This decomposition is easier to discuss, since it only depicts the local state transition.

4.2 Diffusion Approximated Sampling PDF

To approximate adjoint transport solution $L'(\mathbf{x}_{k+1}, -\boldsymbol{\omega}_{k+1}, T_{\text{res},k+1})$, we introduce the transient diffusion equation (DE) into the approximation of transient radiance. The radiant flux solution $\Phi(\mathbf{x}, T)$ of DE in an infinite homogeneous scattering medium satisfying $\sigma_s \gg \sigma_a$ [Contini et al. 1997] is given by :

$$\Phi(\mathbf{x}, T) = \frac{cH(T-\tau)}{\left[4\pi cD(T-\tau)\right]^{3/2}} \exp\left(-\frac{\|\mathbf{x}-\mathbf{x}_e\|^2}{4cD(T-\tau)} - \sigma_a c(T-\tau)\right), \quad (9)$$

where $D = \frac{1}{3(\sigma_a + \sigma_s(1-g))}$,

τ is the time point when the delta wavefront is emitted; \mathbf{x} and T denotes the position and time point to be evaluated, respectively; g is the anisotropic coefficient of Henyey-Greenstein phase function, and is used to compute the reduced scattering coefficient D [Lister et al. 2012]; $H(\cdot)$ is the Heaviside step function which prevents non-causality. This solution can be used to evaluate approximated radiance with specific residual time. Note that Equation (9) doesn't account for the direction information, therefore, the the directed

radiance is approximated based on the radiant flux $\phi(\mathbf{x}, T)$ without direction information.

The above yields part of the distance sampling PDF, and transmittance is accounted for additionally. This would approximate the transient radiance emitted from the emitter, scattered in the medium (approximated transient radiance) and get attenuated from the sampled position to the ray starting position (transmittance). Therefore, the distance sampling PDF takes the following form:

$$p_t(d_k) = \frac{\overbrace{\sigma_e \exp(-\sigma_e d_k)}^{\text{transmittance}} \overbrace{\Phi(\mathbf{x}_k + \boldsymbol{\omega}_k d_k, T_{\text{res},k} - \frac{\eta}{c} d_k)}^{\text{flux approximated transient radiance}}}{Z_k}, \quad (10)$$

Z_k is the normalization constant. Substitute (10) into (7), we have:

$$\tilde{L}_k(\mathbf{x}_k, -\boldsymbol{\omega}_{k+1}, T_{\text{res},k}) = \frac{L'(\mathbf{x}_{k+1}, -\boldsymbol{\omega}_{k+1}, T_{\text{res},k+1})}{\underbrace{\Phi(\mathbf{x}_k + \boldsymbol{\omega}_k d_k, T_{\text{res},k} - \frac{\eta}{c} d_k)}_{\text{variance inducing}}} \frac{\sigma_s Z_k}{\sigma_e}, \quad (11)$$

It can be seen that the denominator of the variance inducing term in Equation (11) is the approximation to the directionally integrated numerator, given the physical meaning of radiant flux. Also, it has been proved that when $\sigma_s \gg \sigma_a$ [Contini et al. 1997], DA can accurately describe the radiant flux distribution and the radiance will be distributed uniformly in terms of direction, thus enabling the flux approximation for radiance. In this case, the variance inducing part can be regarded as a constant. Therefore, the proposed sampling approach achieves much lower variance compared to the existing methods while simultaneously incorporating path time information.

4.3 Sample Generation

Unfortunately, the PDF given by Equation (10) can not be analytically integrated due to its mathematical complexity, which makes analytical inverse sampling infeasible. Therefore, we generate samples according to Equation (10) via resampled importance sampling (RIS) [Talbot 2005]. The sample generation process is as follows:

4.3.1 Scattering event sampling. Since PDF in Equation (10) is actually conditioned on scattering events, we start by sampling the scattering events using exponential sampling. For a given ray, if the maximum travelling distance before hitting an opaque surface is d_m , the probability for medium scattering event can be given by:

$$p_m := P(d_k < d_m) = \int_0^{d_m} \sigma_e \exp(-\sigma_e t) dt = 1 - \exp(-\sigma_e d_m), \quad (12)$$

The sampled scattering event can then be defined by a Bernoulli distribution $\text{Bern}(p_m)$ and can be decided by sampling from this distribution. DA distance sampling will be used if the current scattering event is sampled as a medium event.

4.3.2 Candidate sample generation. For medium scattering events, distance samples can be obtained through Equation (10) via RIS. We use truncated exponential distribution as our candidate distribution, since the local radiance transport is largely influenced by transmittance. Although other sampling methods such as equianangular sampling [Kulla and Fajardo 2012] have been considered as proposals, they are found to be less effective (See supplementary

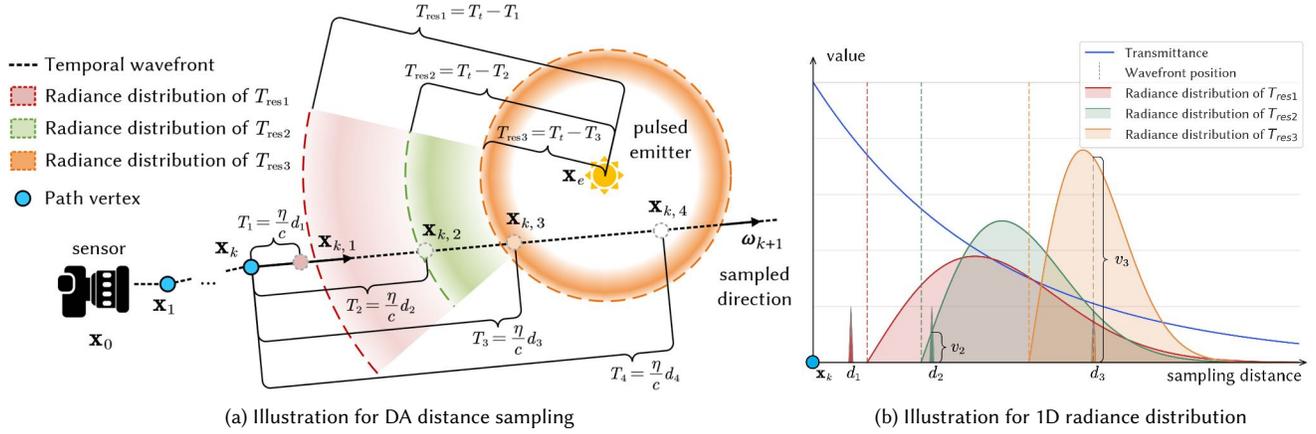


Fig. 4. Illustration of DA distance sampling. Samples are drawn with statistically higher contribution according to product of transmittance and approximated transient radiance: 4 candidate samples d_i are depicted in (a). d_1 and d_4 are invalid. Though d_2 has higher transmittance, the incident contribution on vertex x_k is lower than that of d_3 , due having lower estimated radiance. Note that d_4 is not presented in (b), since it is invalid due to non-causality.

note Figure VII). The candidate sampling distribution in RIS takes the following form:

$$p_{\text{candidate}}(d_k | d_k < d_m) = \frac{\sigma_e \exp(-\sigma_e d_k)}{p_c}, d_k \in [0, d_m), \quad (13)$$

$$d_k = \log(1 - p_c \epsilon) / \sigma_e, \epsilon \sim U[0, 1), p_c = 1 - \exp(-\sigma_e d_m) \quad (14)$$

4.3.3 *Evaluating transient DE and transmittance.* Then, the sampled candidates are used to evaluate the RIS weights. After sampling according to the RIS weights, we can get the following estimator for the transmittance:

$$\tilde{I}(d_k) = \frac{\exp(-\sigma_e d_k)}{\hat{p}(d_k) N_{\text{RIS}}} \sum_{i=1}^{N_{\text{RIS}}} w(d_k), w(d_k) = \frac{\hat{p}(d_k)}{p_{\text{candidate}}(d_k)}, \quad (15)$$

$\hat{p}(d_k)$ is the denominator of Equation 10. The number of candidates N_{RIS} is chosen to be the power of 2, and we use 8 to balance SIMD vectorization efficiency and output quality through out our implementation. To further enhance resampling efficiency, we propose to randomly sample one row from a pre-computed table (32 by 8) of Sobol sequence, then randomly offset each element in the sampled row by a uniformly distributed factor ϵ_0 before the modulo operation into $[0, 1)$ range. Pre-computed Sobol sequence prevents the clustering of candidate samples. It also maintains good randomness and low correlation, while reducing the heavy computational load induced by Sobol sampler state updates. The RIS procedure employed eliminates the resolution-quality trade-off of tabulated sampling while remaining memory-efficient and can be further improved by reservoir resampling [Bitterli et al. 2020].

We provide an intuitive illustration in Figure 4 with 4 candidate distance samples $d_i, i = 1, 2, 3, 4$. Note that $x_{k,i}$ in the figure are all candidate samples, the actual vertex for the k -th scattering event will be resampled from them. Since the sum of T_{res} and the sampling time T_i should be the given target time T_t , a longer sampling distance results in less propagation time (see $x_{k,3}$ and the **tangerine wavefront** of $T_{\text{res}3}$), and vice versa (see $x_{k,1}$ and the **scarlet wavefront** of $T_{\text{res}1}$). The sample with a higher product value of transmittance

and approximated transient radiance is more likely to be resampled. Note that d_1 and d_4 are invalid since $x_{k,1}$ is out of the wavefront range and $x_{k,4}$ results in negative residual time. The weights of both samples are set to zero. In rare cases, if all candidates are invalid, this sampling approach will degrade to naive exponential sampling. The detection of invalid samples will be discussed in Section 5.2.

5 Elliptical Diffusion Approximated Sampling

In addition to free-path distance sampling, efficient transient rendering requires both effective direction sampling and precise path length control strategies. The complete path length control process involves sampling an additional vertex $x_{\text{ell}} \in \mathbb{R}^3$ by first sampling a direction ω from the current vertex x_k , followed by sampling the distance t according to path length constraints. This procedure can be expressed through a conditional probability decomposition:

$$p(x_{\text{ell}}) = p(x_k + \omega t) \rightarrow p(\omega, t | x_k) = p(t | \omega, x_k) p(\omega | x_k), \quad (16)$$

In this section, we introduce a novel direction sampling method that integrates elliptical path length control with diffusion approximation, which refines the probability distribution function $p(\omega | x_k)$ to achieve lower variance. For the conditional distribution $p(t | \omega, x_k)$, we extend ellipsoidal connection method proposed by Pediredla et al. [2019] to volumetric media, which we refer to as *elliptical sampling*. The combination of these two complementary strategies is termed *elliptical diffusion approximated* (EDA) sampling.

5.1 EDA Direction Sampling

The radiant flux in DA employs first-order spherical harmonic approximation and integrates direction away. We then opt for approximation that can capture direction information of transient radiance.

Before starting a new path sample, a time interval will be sampled (deterministic for time-gated rendering) as the interval for target path time, denoted as $[T_{t,m}, T_{t,M})$. The minimum and maximum residual time range are given by $T_{\text{res},m} = T_{t,m} - T_{e,k}$ and $T_{\text{res},M} = T_{t,M} - T_{e,k}$, respectively, where $T_{e,k}$ is the elapsed time

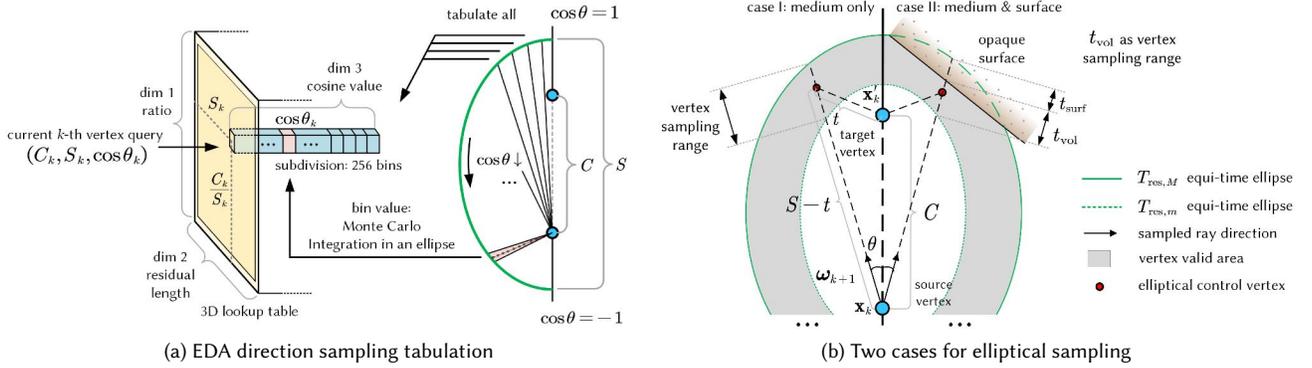


Fig. 5. Two major sampling procedures in EDA sampling. The offline tabulation yields a 3D table for any vertex to query. This table is used for inverse transform sampling. Each value in the table is estimated via Monte Carlo integration on GPU (a). Two possible elliptical sampling cases: the equi-time ellipse does not intersect any surfaces (case I, left half); surface is encountered within the sampling range (case II, right half) (b)

of the path. Importance sampling requires direction sampling to obtain scattering direction that generates radiance samples (1) with high contribution and (2) with the ToF ranging $[T_{res,m}, T_{res,M}]$. Since the upper bound of the residual time at k -th vertex is $T_{res,M}$, for any given direction, the next scattering vertex must reside in an ellipsoid defined by the current vertex, target emitter vertex and residual time, as the next path vertex outside of this ellipsoid will result in longer paths than required. Therefore, we approximate the incident radiance within the ellipse in direction $-\omega$ by the following equation:

$$L_i(\mathbf{x}_k, -\omega, T) = \sigma_s \int_0^{t_M} \exp(-\sigma_e t) L_o(\mathbf{x}_k + \omega t, -\omega, T - \frac{\eta}{c} t) dt \quad (17)$$

$$\approx \sigma_s \int_0^{t_M} \exp(-\sigma_e t) \Phi(\mathbf{x}_k + \omega_k t, T - \frac{\eta}{c} t) dt \quad (18)$$

As shown in Equation (18), L_o can also be approximated by radiant flux introduced in Equation (9). Therefore, if the approximation is sufficiently accurate, we can effectively obtain the volumetric incident radiance from any given direction and for any specific ToF using the convolution of transmittance and DA within an ellipsoid. t_M denotes the polar distance given cT/η as the path length, which can be calculated deterministically:

$$t_M(\cos \theta) = \frac{S^2 - C^2}{2S - 2C \cos \theta}, S = \frac{cT}{\mu}, \quad (19)$$

where C is the focal distance between two vertices being connected; S denotes the length sum of the connection paths; θ is the angle between the sampled direction ω and the major axis vector of the ellipsoid.

Unfortunately, the integral in Equation (18) has no analytical form. In order to draw direction samples from approximated incident radiance, we tabulate Equation (18) by a 3D table. As shown in Figure 5 (a), the first two dimensions, C/S and S are tabulated for conditioning, since C (distance to target) and S (residual time) determine the shape of the ellipsoid, and are known for the given vertex. The resulting direction sampling PDF is conditioned on the

two parameters that will be online-queried. We use C/S instead of C for the first dimension since the former is bounded in $[0, 1]$. The third dimension is the angular dimension used for inverse transform sampling. Note that $\cos \theta$ sampled by our method is the cosine value for the angle between sampled direction and ellipsoid major axis. Interval $[-1, 1]$ is uniformly subdivided into 256 bins, and for each bin, we evaluate Equation (18) via Monte Carlo integration. The ϕ angle is considered isotropic and sampled uniformly in $[-\pi, \pi]$ on the tangent plane defined by the major axis. Therefore, tabulation is evaluated in a 2D ellipse instead of a 3D ellipsoid. Thus we refer to this method as *elliptical DA* instead of *ellipsoidal*.

The tabulation is calculated offline and parallel computation is straightforward. Our offline tabulation only takes around 5 seconds on a single consumer-end GPU and therefore the time overhead is negligible (refer to Section B.3 in supplementary note for more detail). Since we choose to sample $\cos \theta$ and ϕ , the measure of the sampling PDF is mathematically equivalent to solid angle measure, and can thus be directly combined with phase function sampling via one-sample-model MIS [Veach 1998] with balance heuristic. In our implementation, we choose the following adaptive parameter to decide when to use the proposed sampling method:

$$\gamma = \frac{C/S}{C/S + \alpha} \in [0, 1], \alpha \geq 0, \quad (20)$$

where α is a parameter that controls the preference over proposed sampling, and we usually choose 0.5 for experiment. γ defines the probability of choosing the proposed method, and it is adaptive to the shape of the ellipse: when the ellipse resembles a circle (as C/S is close to 0, when bounce count is low), phase function sampling is preferred; as the ellipse flattens after simulating multiple scattering, the proposed sampling is preferred.

To our knowledge, existing work does not importance-sample transient radiance contributions for paths of a given length. In contrast, our DA distance sampling and EDA direction sampling methods effectively address this challenge.

5.2 Elliptical Sampling

For path length control, we adapt the ellipsoidal connection method proposed by Pediredla et al. [2019] to sample a control vertex in participating media. However, their reparameterization method, which is based on the polygon around the center of the (projected) ellipsoid, cannot be directly applied in participating media. This approach also hinders importance sampling and direction reuse, as the PDF of the control vertex \mathbf{x}_{ell} cannot be represented by Equation (16). To address these problems, We introduce a new parameterization that enables efficient sampling in scattering media and supports both importance sampling and direction reuse.

We propose parameterizing the sampling space around the current path vertex (focal point) in the polar coordinate system, which enables the decomposition in Equation (16). This representation simplifies sampling difficulty by breaking it down into the product of two conditional probabilities: one defined by a 1D distance measure and the other by a solid angle measure. In this case, \mathbf{x}_{ell} (the red dot in Figure 5 (b)) can be sampled from a 2D elliptical ring in the media, given the residual time range $[T_{\text{res},m}, T_{\text{res},M}]$. First, a direction ω is sampled using the method in Section 5.1, and then a polar distance t is sampled. Note that the scattering direction ω generated by EDA sampling is reused for this process. Consequently, the next path vertex \mathbf{x}_{k+1} , the sampled vertex \mathbf{x}_{ell} and the current vertex \mathbf{x}_k lie on the same line. This alignment allows us to reuse ray-scene intersection results, significantly reducing rendering time in scenes with high triangle counts.

The sampling of distance t entails two primary scenarios, depending on whether the surface interaction is considered. In the case where the elliptical ring does not intersect any surface in the given direction, the t sample depicted in the left half (case I) in Figure 5 (b), can be generated by the following equation:

$$t = \frac{S^2 - C^2}{2S - 2C \cos \theta}, S \sim P[S_m, S_M], S_m = \frac{cT_{\text{res},m}}{\eta}, S_M = \frac{cT_{\text{res},M}}{\eta} \quad (21)$$

Truncated exponential distribution is used for sampling distribution P , which truncates exponential distribution in $[0, S_M - S_m]$. Note that if uniform distribution is used for P , this sampling method will degrade to one of the sampling method proposed by Jarabo et al. [2014], yet their purpose is to uniformly distribute path vertices. However, since the transmittance for the connection path is $\exp(-\sigma_e S)$, the truncated exponential distribution favors paths with higher overall transmittance, leading to more samples with lower S values. This aligns with our goal of importance sampling based on overall radiance contribution. The sampling PDF takes the form as Equation (22), and full derivation can be found in supplementary note (Section A.2).

$$p(t) = p(S) \left(\frac{dt}{dS} \right)^{-1} = \frac{\eta(S-t)\sigma_e \exp(-\sigma_e(S-S_m))}{c(S-C \cos \theta)(1 - \exp(-\sigma_e(S_M - S_m)))}, \quad (22)$$

The other case involves surface events within sampling range, as shown in the right half (case II) Figure 5 (b). The detection for this case is performed by comparing the closest surface distance t_s with the polar distance t_M concerning the ellipse defined by S_M . Since sampling beyond surfaces is not feasible, we first sample scattering events as in Section 4.3, with $p_m = t_{\text{vol}} / (t_{\text{vol}} + t_{\text{surf}})$. For

a surface event, we directly output t_s , while for a medium event, elliptical sampling is employed and the upper bound for S is updated according to t_s .

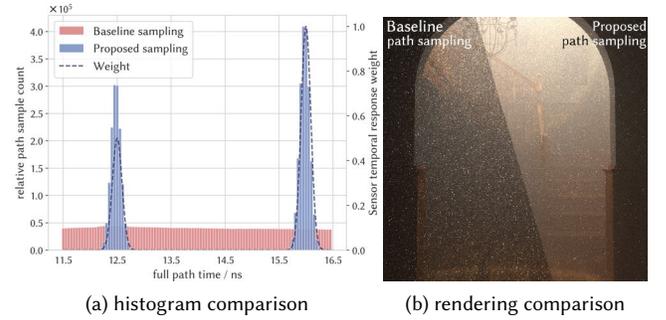


Fig. 6. Importance sampling temporal response function. We record full path samples in a temporal histogram during rendering. Histogram (a) shows that the **baseline sampling** method struggles to generate path samples in time intervals with higher sensor response weights, in contrast to the **proposed sampling method**. Consequently, the baseline sampling approach generates a substantial number of path samples with little overall contribution, resulting in considerably noisier rendered output (b, left half). The temporal response weight comprises two peaks, and two brighter rings are noticeable in the rendered image (b), one on the floor and the other on the wall.

We summarize the abilities of our elliptical sampling method: first, it is able to get the upper bound for distance samples for any residual time via Equation (21), which helps to identify non-causal samples mentioned in Section 4.3; second, the upper bound can also be used for the early identification of paths exceeding the target time range, leading to earlier exit for the path construction and further accelerate rendering by around 1.5 times. Also, its path length control ability can prevent sample rejection caused by $W(\|\bar{\mathbf{x}}\|)$, given the time interval to be sampled from. We present an example of importance-sampling a known $W(\|\bar{\mathbf{x}}\|)$ with two peaks. The results presented in Figure 6 verify the effectiveness of the proposed method: the two peaks of $W(\|\bar{\mathbf{x}}\|)$ are depicted with dashed line in Figure 6(a). The baseline method refers to the direct ToF extension of steady state method. Since the baseline sampling method relies on direct shadow connection, it fails to generate samples falling within the regions with high weights, resulting in a significant portion of samples (85.31%) being rejected due to zero weight. In contrast, our method is able to capture the shape of the response weight and thus avoids sample rejection (<3%). The rendering output in Figure 6(b), which contains two bright rings, shows that our method significantly improves rendering quality. We also provide a proof regarding the optimality of the proposed path control in our supplementary note (Section A.3).

6 Results and Evaluation

We compare the proposed method with path tracing (PT) [Jarabo et al. 2014] and the photon based methods [Liu et al. 2022; Marco et al. 2019], which are the state-of-the-art methods in ToF rendering. All the photon based methods being compared are progressive. We

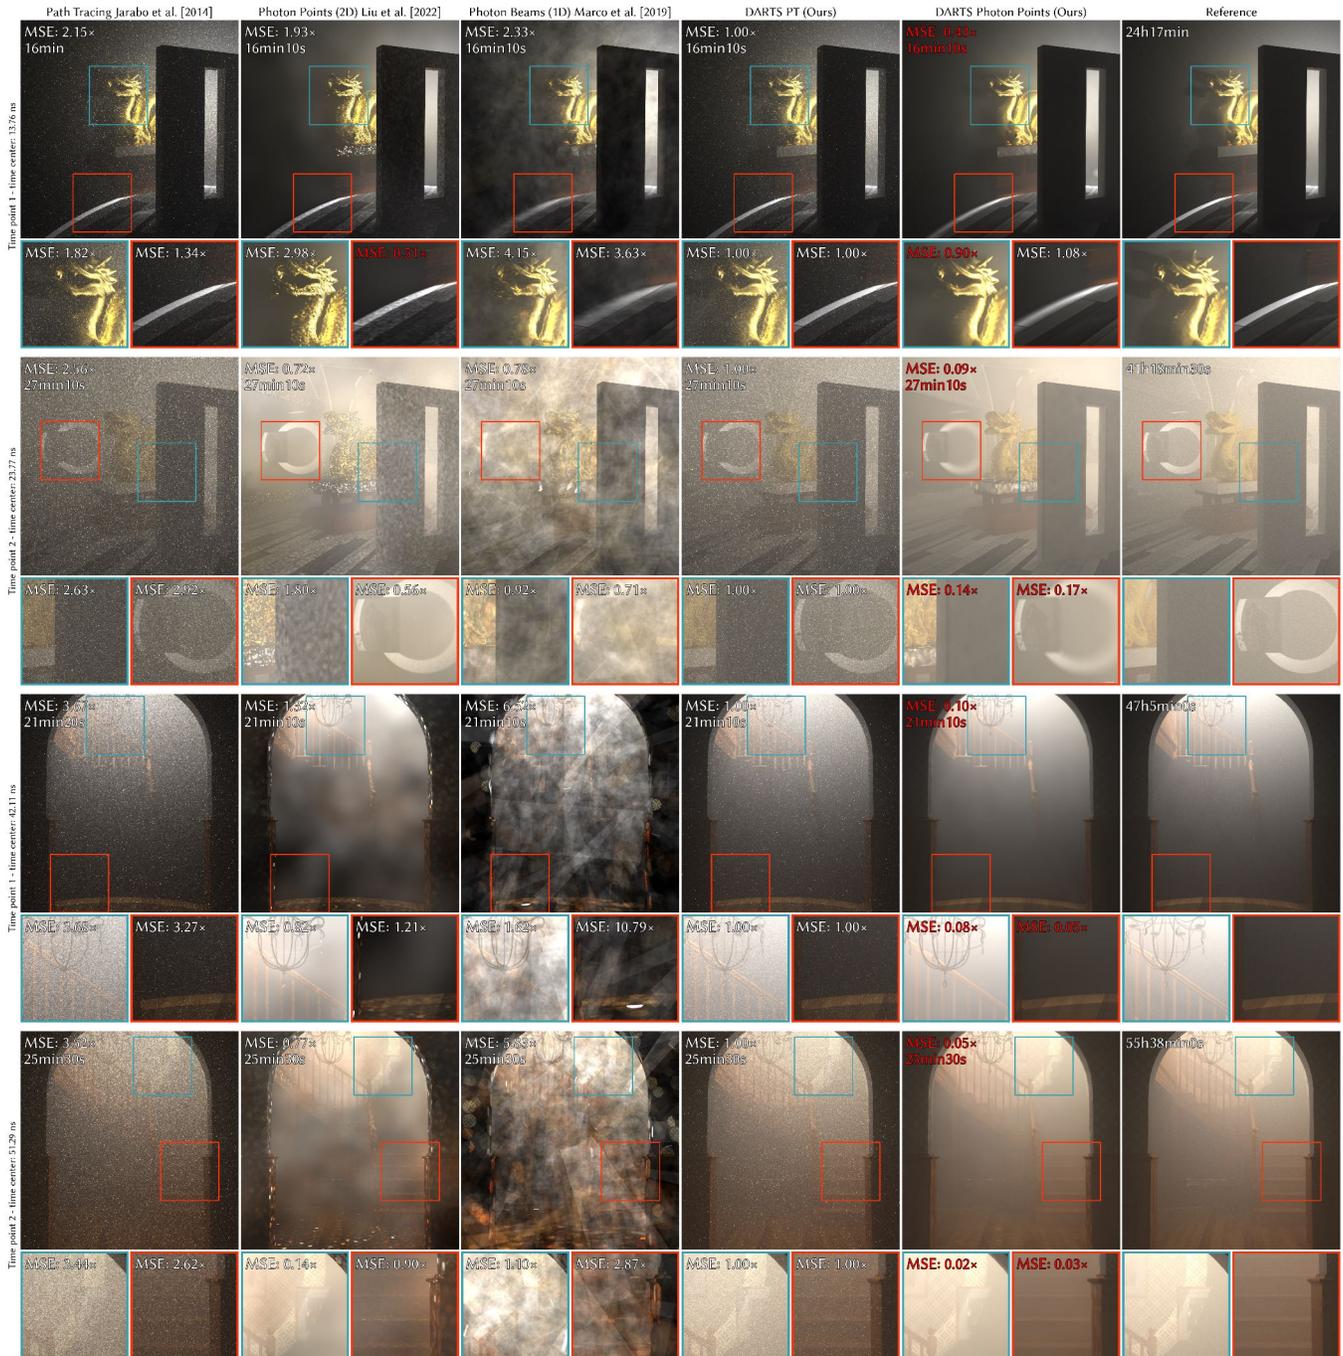


Fig. 7. Time-gated rendering experiments. For each scene, we present the time-gated images rendered by five different methods (organized by column) in two different time intervals (shorter and longer target path lengths). Each image will have two selected areas displayed in magnified views. For original and the proposed methods, we use 6k SPP to render each image and the photon based methods have rendering time equal to the proposed method. We set the MSE of our DARTS PT to be 1, and MSE relative ratios and rendering time are displayed on the top right of each image. The best statistics (row-wise) are highlighted in red. The leftmost texts describe the time point and the corresponding interval center of the rendered images. The original outputs are in HDR format, we therefore normalize the images with their 0.99 quantiles and clamp the output to $[0, 1]$. It can be seen that our DARTS PT (4th column) and DARTS photon points (5th column) methods have greatly improved rendering quality. Notably, photon points equipped with DARTS can achieve noise-free rendering with no obvious visual artifacts.

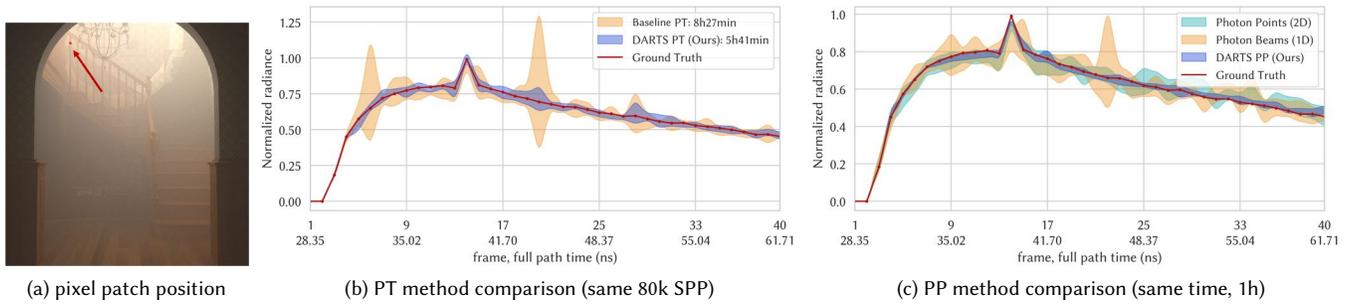


Fig. 8. Transient rendering experiments. We present the transient curve produced by taking the average of the 6×6 pixel patch in (a). The error distribution curves are smoothed with spline interpolation. Equal-SPP comparison between the baseline renderer and our DARTS based one is given in (b), with the rendering time given in the graph legend. Equal-time (1h) comparison between baseline photon based methods and our DARTS based photon point is given in (c). Both curves show that DARTS is able to reduce the asymptotic variance and improve the rendering quality.

Table 1. Properties of the test scenes. Note that scattering intensity is measured by the camera-to-emitter transport mean-free-path (TMFP) distance. Higher value means more expected scattering events along the distance.

Properties \ Scenes	GLOSSY DRAGON	STAIRCASE [Bitterli 2016b]
Triangle Count	87k	262k
PT rendering time	Equal to that of DARTS PT	Equal to that of DARTS PT
Scattering intensity*	3.2 TMFP	4.3 TMFP
Surface material	glossy / specular reflection	plastic coated material

present time-gated rendering and transient rendering experiments in Section 6.1 and 6.2, respectively. Additionally, we provide numerical results in Section 6.3 to demonstrate the stability of our algorithm under various scene and sensor settings. We implement our sampling method in the unidirectional PT of *pbrt-v3* [Pharr et al. 2023], photon points and photon beams (PP and PB for short, respectively) methods of *Tungsten* [Bitterli 2018]. These renderers are verified to converge to the same results with correct settings and code adjustments (Section B.1 of our supplementary note). Unlike other works that mostly focus on volumetric transport and low order scattering, our experiment setup enables full transport and great maximum allowed depth (such as 200 bounces). All the experiments are done on 112-core Intel Xeon Platinum 8280 CPU@2.60GHz with 104 threads. GPU used for tabulation is Titan RTX. We also run the rendering to converged state to ensure that DARTS remains unbiased and consistent. All rendered HDR images are in the linear RGB color space unless stated otherwise. The reference images are produced by path tracing method.

6.1 Time-gated Rendering

In Figure 7, we compare our DARTS based renderer with: (1) A unidirectional path tracing version of the methods proposed by Jarabo et al. [2014], which is referred to as *uniform-time* PT below. (2) 1D progressive transient photon beam method [Marco et al. 2019] reproduced by Liu et al. [2022] (3) Progressive photon point (point-beam 2D, [Křivánek et al. 2014]) implemented by Liu et al. [2022]. The proposed path sampling method is tested both on PT (denoted by DARTS PT) and photon based method (DARTS photon

points, PP for short). For each scene, two time intervals (shorter and longer target path lengths) are rendered and presented as the first and the second row of the related scene, respectively. Note that the estimators proposed by Liu et al. [2022] produce results inferior to those listed above in camera-warped and full transport settings, and some of them needs mathematical corrections to work under camera-warped settings, which is out of this paper’s scope and are not included (see Figure X in the supplementary note). The methods proposed by Jarabo et al. [2014] is modified to fit the unidirectional path tracing (see Section B.4 in our supplementary note for more detail). The presented images are normalized with their 0.99 quantile numbers. In the following, we mainly provide two scenes for comparison. More scenes and results can be found in our supplementary note, interactive local web-viewer and video.

Some of the important properties of the two test scenes: GLOSSY DRAGON (first two rows of Figure 7) and STAIRCASE [Bitterli 2016a] (last two rows of Figure 7) are listed in Table 1. Both volumetric scenes feature complex geometries and surface transport properties. The pulsed point emitters in each scene are not within the direct line-of-sight. We render the scene with 6k SPP for DARTS PT, and the compared methods employ the same rendering time as DARTS PT. Quantitative results presented in Figure 7 are normalized by the MSE of DARTS PT results.

It is evident that our DARTS PT significantly outperforms the *uniform-time* PT in output quality, reducing MSE metrics by about 2-4 times. Moreover, DARTS PT can already outperform photon based methods in most scenes due to its bias-free nature, while DARTS PP further increased the gap between our methods and the compared methods. We also observe that, though progressive method is employed, rendering under a fixed time budget and accounting for full transport introduces bias that proves challenging to mitigate, even with extensive parameter tuning. This challenge is particularly pronounced in the case of the photon beams method, leading to noticeable visual artifacts and a performance drop. With our sampling method, photon based methods can adopt lower photon counts per sample to achieve higher SPP and the parameters are easier to tune, and thus have better convergence.

6.2 Transient Rendering

Our method is better suited for time-gated rendering where temporal path reuse [Jarabo et al. 2014] isn't employed, but we still report performance improvements in transient rendering using our proposed methods. Similar to time-gated rendering, we compare PT and PP with our sampling method against their baseline counterparts and baseline photon beams. Here we present comparison experiment for STAIRCASE scene. Each frame has the same temporal width as time-gated experiments and 40 frames instead of only 1 frame are rendered. Since it is generally challenging for PT to produce transient images with low variance, here we compare path tracing methods (rendered with 80k SPP) and photon based method (rendered for the same time, 1h) differently. The transient curve from a 6×6 pixel patch shown in Figure 8(a) is presented. The error distribution of the PT method is presented in Figure 8(b), whereas the error distribution of the photon-based method is depicted in Figure 8(c).

The results presented in Figure 8 show that compared to baseline methods, the proposed method is able to improve the rendering quality with significantly less rendering time. However, photon based methods outperform path tracing based renderer with the proposed sampling methods, due to the utilization of temporal path reuse and spatial-temporal blurring (rectangle kernel for temporal blurring). We further test the photon points methods with our proposed sampling methods, observing improvements in both rendering quality and rendering speed.

The reduction in rendering time is achieved by eliminating temporal path reuse: the proposed method constructs paths and connections for a specific frame, and the subsequent random walk are terminated once the path time exceeds the duration of that specified frame, thus requiring fewer overall bounces. The rendering quality is improved due to the ability of importance-sampling the whole path according to transient radiance contribution. Besides, the proposed methods are able to uniformly distribute the path samples in time-domain, similar to the work of Jarabo et al.'s [2014] but in a different way: our work is able to place equal number of path samples in each frame (statistically, regardless of visibility term) and our constructed path is not reused across frames but dedicated to a single frame instead. Thus, our path samples do not need to compromise among all the frames.

6.3 Parametric Variation and Further Comparisons

In the following, we mainly discuss the performance of our method under different scattering coefficient σ_s , temporal gate width and total rendering time. This section concludes with an ablation study of our sampling method in the CORNELL BOX scene. Unless otherwise specified, all subsequent experiments employ 2k SPP for PT, while photon-based methods adhere to the same rendering time as DARTS PT. Rendered figures are available in the supplementary note and our interactive local web-viewer.

6.3.1 Scattering coefficient. For both GLOSSY DRAGON and STAIRCASE scene, we employ six different σ_s settings, ranging from low-order to high-order scattering. In each scene, we select the same two time intervals as in the time-gated rendering experiments for rendering, and 30 images are rendered to calculate MSE. For the sake

of clarity, we convert the scattering coefficient to TMFP values for path length of the corresponding time interval. To obtain reasonable outputs, all photon-based methods require parameter tuning for different scattering coefficients. In the following figures, we present the comparison results in two above-mentioned scenes. It can be seen from Figure 9 that our DARTS PT maintains lower variance compared to both baseline PT and photon-based methods, exhibiting greater stability when the scattering coefficient varies. Meanwhile, DARTS PP gradually gains its superiority as the scattering media get denser but in thin scattering media, its performance is inferior to path tracing methods, due to the surface transport. Compared with the experiments presented in Section 6.1, photon-based methods employ less rendering time, leading to fewer SPP, which can impact progressive rendering and result in a performance drop.

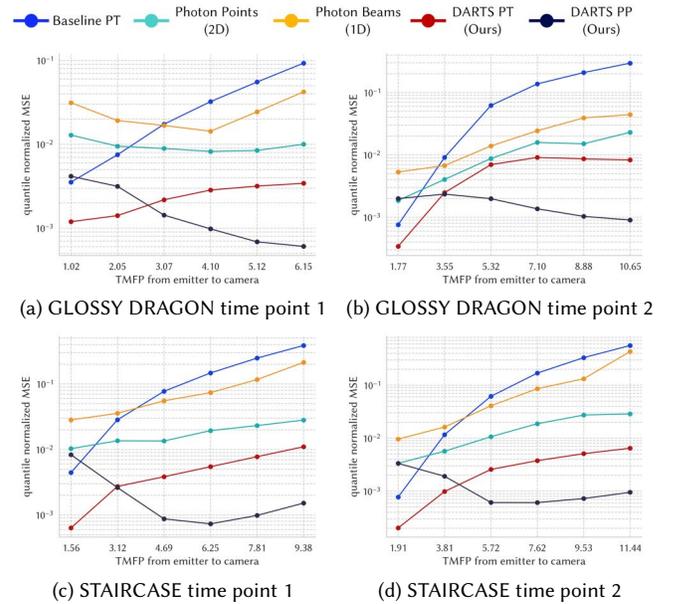


Fig. 9. Scattering coefficient & MSE relationship curves. Curves depicting the relationship between scattering coefficient (converted to TMFP) and MSE are presented. The first and the second row are obtained from GLOSSY DRAGON and STAIRCASE scene, respectively. Note that multiple scattering at higher σ_s values contributes to an overall decrease in signal amplitude and, consequently, a reduction in MSE. To account for this, we normalize the rendered results using the 0.95 quantile before MSE calculation.

6.3.2 Temporal gate width. Here we present the gate width experiments in the STAIRCASE scene. Six gate widths are selected, ranging from a narrow delta-function-like gate to a wide steady-state-like gate. Likewise, the gate width ΔT is converted into a ratio with the mean free path (MFP) of the scene. The scene settings align with those of time-gated experiments described in Section 6.1. The curves are given in Figure 10. We can observe that for wider time gates, as the setup more closely resembles steady-state rendering, the advantages of DARTS PT in time-resolved sampling are less observable over path tracing methods. Therefore, the gap between the DARTS PT and photon points methods is gradually narrowing.

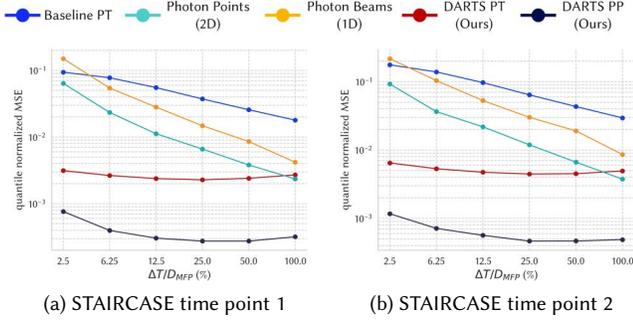


Fig. 10. Temporal gate width & MSE relationship curves. The test scene is STAIRCASE scene and the outputs are also normalized with 0.95 quantile. Our proposed method is seen to be stable across different settings.

6.3.3 Numerical convergence. A straightforward numerical convergence analysis is depicted in Figure 11, where different SPP settings (scaled by a factor of ten) are employed to render the first time point of both GLOSSY DRAGON and STAIRCASE scenes. For photon based methods, we have rendered the scenes with rendering time equivalent to DARTS PT. The result shows that: for PT methods, the proposed path sampling method is able to improve the MSE convergence by at least an order of magnitude in this scene; for photon based method, the convergence for fewer-sample cases are significantly improved, whereas as the SPP increases, the improvement is less significant and the performance can even be inferior to DARTS PT. This is caused by the inherent bias of photon based methods, which is difficult to eliminate, even though the methods are progressive.

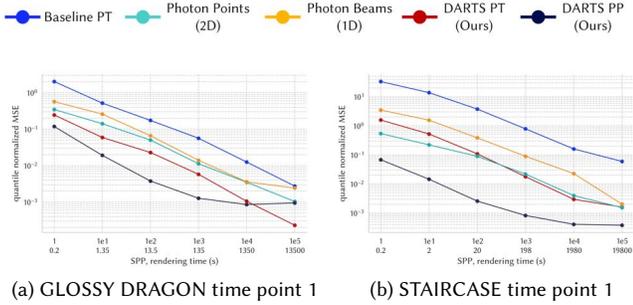


Fig. 11. Numerical convergence curve. We compare the proposed method with baseline method and photon based methods under several different SPP (thus, rendering time) settings in both scenes. Both SPP (for path tracing methods) and rendering time (for both path tracing and photon based methods) are presented in the x-axis labels. Note that baseline PT uses the same number of SPP as DARTS PT, and the baseline PT takes approximately 1.5 times longer to render compared to DARTS PT.

6.3.4 Ablation study. Our proposed sampling method consists of two parts: DA distance sampling and EDA sampling. We validate the ability of these parts in improving rendering quality through several strategies: baseline renderer, the renderer with only partial part or one of the proposed methods and the renderer with both methods

enabled. Here we use a modified version of the CORNELL BOX scene with a pulsed point source. The scene is rendered under various scene settings such as different scattering coefficients and temporal gate widths. Both qualitative and quantitative results under one example setting are presented in Figure 12. It is evident that the collaborative utilization of proposed methods substantially improves rendering quality. For additional results of other settings, one can refer to our supplementary materials (Figure III in Section C.1).

7 Discussion and Limitations

In summary, we have proposed a novel path sampling method to improve the overall quality of ToF rendering tasks. The proposed method is derived based on diffusion approximation in homogeneous scattering media and its Monte Carlo integration in a residual-time-defined ellipse, together with the non-trivial extension of ellipsoidal connection that can be directly applied for scattering media. Our experiments demonstrate that the proposed method is able to improve rendering quality and efficiency in both path tracing and photon-based methods. Moreover, the improved path tracing method performs comparably to, and in some cases surpasses, photon-based methods in scenes with scattering media and complex surface properties. We anticipate that our work will contribute to the simulation of increasingly sophisticated ToF sensors in the field of optics and sensing.

In the following, we first discuss the implementation differences of the proposed method in different frameworks, and then provide the limitations and the future avenues for this work.

Implementation differences in different frameworks. The previous derivations mainly stem from PT framework. While the same methodology can be applied to photon-based method, we do highlight two major differences in implementation: (1) DA distance sampling and EDA sampling are only used in photon pass, and the camera will be considered as a virtual importon [Vorba et al. 2014] emitter, while in PT, since the path starts from the camera, the real emitter is used. (2) In photon based methods, we cache the control vertex in the photon map and leave the gathering for the sensor pass, while in PT, a generalized shadow connection will be made to evaluate the radiance immediately after sampling a control vertex. For detailed explanation of the differences in implementation, one can refer to Section B.2 in our supplementary note.

7.1 Limitations & Future Avenues

7.1.1 Emitter & phase function types. The emitter in our theoretical derivation is assumed to be a point emitter. We employ two types of emitters in the experiments: point emitters and spot emitters. The latter is employed in photon-based methods, where the camera acts as a spot emitter of importons. For more complex emitters, for example, collimated emitters can be handled by sampling single scattering events and treating each scattering vertex as a virtual point emitter, while the area emitters can be approached by point sampling, with each sample considered as an equivalent spot emitter. The phase function employed is currently assumed to be Henyey-Greenstein phase function, yet we note that our method might be extended to phase functions that have diffusion approximation, such

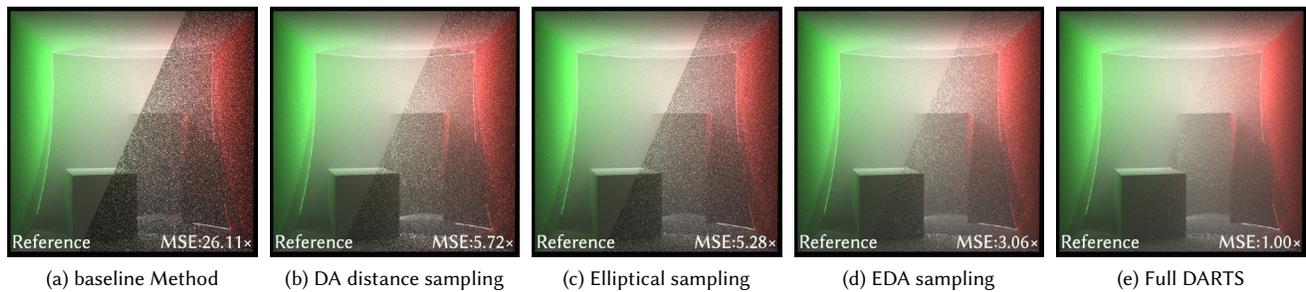


Fig. 12. Ablation study: an example case. Comparison between DARTS and four alternative strategies is presented under the same rendering time settings. The images shown above are synthesized by averaging 40 images (totaling 30 min for each case), with the reference being rendered for 4.5h. Images shown are normalized by their 0.99 quantile. We normalize all the MSE results by the MSE of images produced by DARTS. In each image, the left half is the reference and the right half employs strategies as indicated by the sub-captions. The proposed method results in a substantial reduction in MSE (in this case, variance).

as microflake phase function [Jakob et al. 2010]. This part is left for future work.

7.1.2 Violation to the assumption of DA. Our current work addresses the ToF rendering challenges in homogeneous scattering media where the scattering exhibits limited directionality. Media with heterogeneity, strong absorption effect and strong directionality may violate the assumptions of diffusion approximation. In such cases, diffusion theory may not be applicable when incorporating global transient radiance information. Experiments on the effects of absorption coefficients and phase directionality are conducted, and fortunately, our method can retain its robustness (see supplementary note Section C.5). Additionally, as elliptical sampling in EDA involves a two-step approach, the reused direction from EDA may be sub-optimal for peaky phase function theoretically. However, we find that the 2D radiance contribution of the product of two consecutive phase functions for elliptical connection does not differ much from the single phase function case (see supplementary note Section C.6), therefore, local phase function still works well for MIS. For scenes with homogeneous scattering media in distinct regions, our work can still be applied in bidirectional methods, where vertices on emitter paths can be regarded as point sources.

7.1.3 Scenes with complex visibility. Using diffusion theory in scenes with complex visibility lacks robust theoretical substantiation. However, for efficiency, visibility in RIS can be overlooked, as noted by Bitterli et al. [2020]. Our experiments indicate that, even when emitters are not directly visible, our sampling method can still greatly enhance rendering quality. However, incorporating visibility term in unidirectional renderers poses inherent challenges. For example, calculating ray-scene intersection for every candidate sample during RIS becomes computationally burdensome. This overhead might outweigh the performance gains, even if strategies are utilized for acceleration, like caching a shadow map for the emitter.

For future works, it would be interesting to investigate how our method can be applied for differentiable rendering in transient state, since higher sampling efficiency generally yields better backward performance. Additionally, investigating the compatibility of our method with more sophisticated rendering frameworks, such as bidirectional path tracing and metropolis light transport, could prove

valuable, particularly in rendering challenging scenes like non-line-of-sight simulation setups. We believe that our study provides inspiration for future researches in time-of-flight sensor simulation, especially in scenarios involving diverse and complex scattering. The testing scenes and code of our method in both *pbrt-v3* and *Tungsten* frameworks can be found in our supplementary material.

Acknowledgments

The work was supported in part by the National Natural Science Foundation of China (62131011), Shenzhen Science and Technology Project under Grant (JSGG20220831095602005), Natural Science Foundation of Guangdong Province, China, (2023A1515012716).

References

- Marco Ament, Christoph Bergmann, and Daniel Weiskopf. 2014. Refractive radiative transfer equation. *ACM Transactions on Graphics (TOG)* 33, 2 (2014), 1–22.
- Benjamin Attal, Eliot Laidlaw, Aaron Gokaslan, Changil Kim, Christian Richardt, James Tompkin, and Matthew O’Toole. 2021. Törf: Time-of-flight radiance fields for dynamic scene view synthesis. *Advances in neural information processing systems* 34 (2021), 26289–26301.
- Benedikt Bitterli. 2016a. Rendering resources. <https://benedikt-bitterli.me/resources/>.
- Benedikt Bitterli. 2016b. Virtual Femto Photography. <https://benedikt-bitterli.me/femto.html>.
- Benedikt Bitterli. 2018. Tungsten Renderer. <https://github.com/tunabrain/tungsten/>.
- Benedikt Bitterli and Wojciech Jarosz. 2017. Beyond points and beams: Higher-dimensional photon samples for volumetric light transport. *ACM Transactions on Graphics (TOG)* 36, 4 (2017), 1–12.
- Benedikt Bitterli, Chris Wyman, Matt Pharr, Peter Shirley, Aaron Lefohn, and Wojciech Jarosz. 2020. Spatiotemporal reservoir resampling for real-time ray tracing with dynamic direct lighting. *ACM Transactions on Graphics (TOG)* 39, 4 (2020), 148–1.
- Daniele Contini, Fabrizio Martelli, and Giovanni Zaccanti. 1997. Photon migration through a turbid slab described by a model based on diffusion approximation. I. Theory. *Applied optics* 36, 19 (1997), 4587–4599.
- Xi Deng, Shaojie Jiao, Benedikt Bitterli, and Wojciech Jarosz. 2019. Photon surfaces for robust, unbiased volumetric density estimation. *ACM Transactions on Graphics* 38, 4 (2019).
- Dongyu Du, Xin Jin, Rujia Deng, Jinshi Kang, Hongkun Cao, Yihui Fan, Zhiheng Li, Haoqian Wang, Xiangyang Ji, and Jingyan Song. 2022. A boundary migration model for imaging within volumetric scattering media. *Nature Communications* 13, 1 (2022), 3234.
- Daniele Faccio, Andreas Velten, and Gordon Wetzstein. 2020. Non-line-of-sight imaging. *Nature Reviews Physics* 2, 6 (2020), 318–327.
- Iliyan Georgiev, Jaroslav Krivanek, Toshiya Hachisuka, Derek Nowrouzezahrai, and Wojciech Jarosz. 2013. Joint importance sampling of low-order volumetric scattering. *ACM Trans. Graph.* 32, 6 (2013), 164–1.
- Tobias Gruber, Frank Julca-Aguilar, Mario Bijelic, and Felix Heide. 2019. Gated2depth: Real-time dense lidar from gated images. In *Proceedings of the IEEE/CVF International Conference on Computer Vision*. 1506–1516.

- Toshiya Hachisuka and Henrik Wann Jensen. 2009. Stochastic progressive photon mapping. In *ACM SIGGRAPH Asia 2009 papers*. 1–8.
- Abderrahim Halimi, Aurora Maccarone, Robert A. Lamb, Gerald S. Buller, and Stephen McLaughlin. 2021. Robust and Guided Bayesian Reconstruction of Single-Photon 3D Lidar Data: Application to Multispectral and Underwater Imaging. *IEEE Transactions on Computational Imaging* 7 (2021), 961–974. <https://doi.org/10.1109/TCI.2021.3111572>
- Sebastian Herholz, Oskar Elek, Jiří Vorba, Hendrik Lensch, and Jaroslav Krivánek. 2016. Product importance sampling for light transport path guiding. In *Computer Graphics Forum*, Vol. 35. Wiley Online Library, 67–77.
- Sebastian Herholz, Yangyang Zhao, Oskar Elek, Derek Nowrouzezahrai, Hendrik PA Lensch, and Jaroslav Krivánek. 2019. Volume path guiding based on zero-variance random walk theory. *ACM Transactions on Graphics (TOG)* 38, 3 (2019), 1–19.
- J Eduard Hoogenboom. 2008. Zero-variance Monte Carlo schemes revisited. *Nuclear science and Engineering* 160, 1 (2008), 1–22.
- Julian Iseringhausen and Matthias B Hullin. 2020. Non-line-of-sight reconstruction using efficient transient rendering. *ACM Transactions on Graphics (ToG)* 39, 1 (2020), 1–14.
- Wenzel Jakob, Adam Arbree, Jonathan T Moon, Kavita Bala, and Steve Marschner. 2010. A radiative transfer framework for rendering materials with anisotropic structure. In *ACM SIGGRAPH 2010 papers*. 1–13.
- Adrian Jarabo. 2012. Femto-photography: Visualizing light in motion. *Universidad de Zaragoza* (2012).
- Adrian Jarabo and Victor Arellano. 2018. Bidirectional rendering of vector light transport. In *Computer Graphics Forum*, Vol. 37. Wiley Online Library, 96–105.
- Adrian Jarabo, Julio Marco, Adolfo Muñoz, Raul Buisan, Wojciech Jarosz, and Diego Gutierrez. 2014. A framework for transient rendering. *ACM Transactions on Graphics (ToG)* 33, 6 (2014), 1–10.
- Adrian Jarabo, Belen Masia, Julio Marco, and Diego Gutierrez. 2017. Recent advances in transient imaging: A computer graphics and vision perspective. *Visual Informatics* 1, 1 (2017), 65–79.
- Wojciech Jarosz, Derek Nowrouzezahrai, Iman Sadeghi, and Henrik Wann Jensen. 2011. A comprehensive theory of volumetric radiance estimation using photon points and beams. *ACM transactions on graphics (TOG)* 30, 1 (2011), 1–19.
- Henrik Wann Jensen and Per H. Christensen. 1998. Efficient Simulation of Light Transport in Scenes with Participating Media Using Photon Maps. In *Proceedings of the 25th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '98)*. Association for Computing Machinery, New York, NY, USA, 311–320. <https://doi.org/10.1145/280814.280925>
- Daiki Kijima, Takahiro Kushida, Hiromu Kitajima, Kenichiro Tanaka, Hiroyuki Kubo, Takuya Funatomi, and Yasuhiro Mukaigawa. 2021. Time-of-flight imaging in fog using multiple time-gated exposures. *Optics Express* 29, 5 (2021), 6453–6467.
- Jaroslav Krivánek, Iliyan Georgiev, Toshiya Hachisuka, Petr Vévoda, Martin Šik, Derek Nowrouzezahrai, and Wojciech Jarosz. 2014. Unifying points, beams, and paths in volumetric light transport simulation. *ACM Transactions on Graphics (TOG)* 33, 4 (2014), 1–13.
- Christopher Kulla and Marcos Fajardo. 2012. Importance sampling techniques for path tracing in participating media. In *Computer graphics forum*, Vol. 31. Wiley Online Library, 1519–1528.
- Ivan T Lima, Anshul Kalra, and Sherif S Sherif. 2011. Improved importance sampling for Monte Carlo simulation of time-domain optical coherence tomography. *Biomedical optics express* 2, 5 (2011), 1069–1081.
- Daqi Lin, Chris Wyman, and Cem Yuksel. 2021. Fast volume rendering with spatiotemporal reservoir resampling. *ACM Transactions on Graphics (TOG)* 40, 6 (2021), 1–18.
- Tom Lister, Philip A Wright, and Paul H Chappell. 2012. Optical properties of human skin. *Journal of biomedical optics* 17, 9 (2012), 090901–090901.
- Yang Liu, Shaojie Jiao, and Wojciech Jarosz. 2022. Temporally sliced photon primitives for time-of-flight rendering. *Computer Graphics Forum (Proceedings of EGSR)* 41, 4 (July 2022). <https://doi.org/10/jgzq>
- Julio Marco, Ibón Guillén, Wojciech Jarosz, Diego Gutierrez, and Adrian Jarabo. 2019. Progressive transient photon beams. In *Computer graphics forum*, Vol. 38. Wiley Online Library, 19–30.
- Julio Marco, Quercus Hernandez, Adolfo Munoz, Yue Dong, Adrian Jarabo, Min H Kim, Xin Tong, and Diego Gutierrez. 2017a. DeepToF: off-the-shelf real-time correction of multipath interference in time-of-flight imaging. *ACM Transactions on Graphics (ToG)* 36, 6 (2017), 1–12.
- Julio Marco, Wojciech Jarosz, Diego Gutierrez, and Adrian Jarabo. 2017b. Transient photon beams. In *ACM SIGGRAPH 2017 Posters*. 1–2.
- Jan Novák, Iliyan Georgiev, Johannes Hanika, and Wojciech Jarosz. 2018. Monte Carlo methods for volumetric light transport simulation. In *Computer graphics forum*, Vol. 37. Wiley Online Library, 551–576.
- Xian Pan, Victor Arellano, and Adrian Jarabo. 2019. Transient instant radiosity for efficient time-resolved global illumination. *Computers & Graphics* 83 (2019), 107–113.
- Adithya Pediredla, Ashok Veeraraghavan, and Ioannis Gkioulekas. 2019. Ellipsoidal path connections for time-gated rendering. *ACM Transactions on Graphics (TOG)* 38, 4 (2019), 1–12.
- Vijitha Periyasamy and Manojit Pramanik. 2016. Importance sampling-based Monte Carlo simulation of time-domain optical coherence tomography with embedded objects. *Applied optics* 55, 11 (2016), 2921–2929.
- Matt Pharr, Wenzel Jakob, and Greg Humphreys. 2023. *Physically based rendering: From theory to implementation*. MIT Press.
- Joshua Rapp, Charles Saunders, Julián Tachella, John Murray-Bruce, Yoann Altmann, Jean-Yves Tournier, Stephen McLaughlin, Robin MA Dawson, Franco NC Wong, and Vivek K Goyal. 2020. Seeing around corners with edge-resolved transient imaging. *Nature communications* 11, 1 (2020), 5929.
- Zhong Ren, Kun Zhou, Stephen Lin, and Baining Guo. 2008. Gradient-based Interpolation and Sampling for Real-time Rendering of Inhomogeneous, Single-scattering Media. In *Computer Graphics Forum*, Vol. 27. Wiley Online Library, 1945–1953.
- Diego Royo, Talha Sultan, Adolfo Muñoz, Khadijeh Masumnia-Bisheh, Eric Brandt, Diego Gutierrez, Andreas Velten, and Julio Marco. 2023. Virtual Mirrors: Non-Line-of-Sight Imaging Beyond the Third Bounce. *ACM Transactions on Graphics (TOG)* 42, 4 (2023), 1–15.
- Kfir Shem-Tov, Sai Praveen Bangaru, Anat Levin, and Ioannis Gkioulekas. 2020. Towards reflectometry from interreflections. In *2020 IEEE International Conference on Computational Photography (ICCP)*. IEEE, 1–12.
- Justin F Talbot. 2005. *Importance resampling for global illumination*. Brigham Young University.
- Eric Veach. 1998. *Robust Monte Carlo methods for light transport simulation*. Stanford University.
- Andreas Velten, Di Wu, Adrian Jarabo, Belen Masia, Christopher Barsi, Chinmaya Joshi, Everett Lawson, Mounji Bawendi, Diego Gutierrez, and Ramesh Raskar. 2013. Femto-photography: capturing and visualizing the propagation of light. *ACM Transactions on Graphics (ToG)* 32, 4 (2013), 1–8.
- Jiří Vorba, Ondřej Karlík, Martin Šik, Tobias Ritschel, and Jaroslav Krivánek. 2014. On-line learning of parametric mixture models for light transport simulation. *ACM Transactions on Graphics (TOG)* 33, 4 (2014), 1–11.
- Amanpreet Walia, Stefanie Walz, Mario Bijelic, Fahim Mannan, Frank Julca-Aguilar, Michael Langer, Werner Ritter, and Felix Heide. 2022. Gated2gated: Self-supervised depth estimation from gated images. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition*. 2811–2821.
- Rihui Wu, Adrian Jarabo, Jinli Suo, Feng Dai, Yongdong Zhang, Qionghai Dai, and Diego Gutierrez. 2018. Adaptive polarization-difference transient imaging for depth estimation in scattering media. *Optics Letters* 43, 6 (2018), 1299–1302.
- Shumian Xin, Sotiris Nousias, Kiriakos N Kutulakos, Aswin C Sankaranarayanan, Srinivasa G Narasimhan, and Ioannis Gkioulekas. 2019. A theory of Fermat paths for non-line-of-sight shape reconstruction. In *Proceedings of the IEEE/CVF conference on computer vision and pattern recognition*. 6800–6809.
- Shinyoung Yi, Donggun Kim, Kiseok Choi, Adrian Jarabo, Diego Gutierrez, and Min H Kim. 2021. Differentiable transient rendering. *ACM Transactions on Graphics (TOG)* 40, 6 (2021), 1–11.
- Tianyi Zhang, Mel J White, Akshat Dave, Shahabuddin Ghajari, Ankit Raghuram, Alyosha C Molnar, and Ashok Veeraraghavan. 2022. First Arrival Differential LiDAR. In *2022 IEEE International Conference on Computational Photography (ICCP)*. IEEE, 1–12.
- Vytautas Zickus, Ming-Lo Wu, Kazuhiro Morimoto, Valentin Kapitanov, Areeba Fatima, Alex Turpin, Robert Insall, Jamie Whitelaw, Laura Machesky, Claudio Bruschini, et al. 2020. Fluorescence lifetime imaging with a megapixel SPAD camera and neural network lifetime estimation. *Scientific Reports* 10, 1 (2020), 20986.