

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

QIANYUE HE, Tsinghua Shenzhen International Graduate School, China

DONGYU DU, Tsinghua Shenzhen International Graduate School, China

HAITIAN JIANG, Tsinghua Shenzhen International Graduate School, China

XIN JIN*, Tsinghua Shenzhen International Graduate School, China

ACM Reference Format:

Qianyue He, Dongyu Du, Haitian Jiang, and Xin Jin. 2024. Supplementary Note: DARTS: Diffusion Approximated Residual Time Sampling for Time-of-flight Rendering in Homogeneous Scattering Media. 1, 1 (September 2024), 20 pages. <https://doi.org/10.1145/nnnnnnnn.nnnnnnnn>

OUTLINE

This supplementary material provides further details and results to support the content from the main paper. It is organized as follows.

Supplementary Notes

Section A: Detailed derivations of the proposed methods and some necessary theoretical discussions.

Section B: Implementation details of DARTS within path tracing and photon-based frameworks.

Section C: More experiment results to showcase the effectiveness of DARTS, including:

- Results of the ablation study in the CONRELL BOX scene.
- Results in two other complex scenes (BATHROOM and KITCHEN scene [Bitterli 2016]).
- Extensions of the proposed methods to the photon beam methods.
- Comparisons with equiangular [Kulla and Fajardo 2012] sampling and uniform sampling.
- Comparisons between unidirectional and bidirectional method (bidirectional path tracing) in CORNELL BOX scene.
- Comparisons with other estimators (transient photon planes and photon volumes) proposed by Liu et al. [2022].
- Investigation for violation to the assumptions of DA and peaky phase function cases.

*Corresponding Author

Authors' addresses: Qianyue He, he-qy22@mails.tsinghua.edu.cn, Tsinghua Shenzhen International Graduate School, China; Dongyu Du, dudy19@mails.tsinghua.edu.cn, Tsinghua Shenzhen International Graduate School, China; 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 ACM 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 Association for Computing Machinery.

XXXX-XXXX/2024/9-ART \$15.00

<https://doi.org/10.1145/nnnnnnnn.nnnnnnnn>

Interactive local web-viewer

All the qualitative results (Section 6.1 & 6.2 of the main paper) and rendered images of the curves (Section 6.3 of the main paper), together with some of the results presented in this supplementary material are presented in our interactive local web-viewer.

Supplementary Video

A supplementary video is presented in our supplementary materials, including: (1) time-gated rendering comparisons in the DINING ROOM, STAIRCASE and GLOSSY DRAGON scenes. (2) transient rendering comparisons in the STAIRCASE scene (the same scene as in the main paper, with higher temporal resolution).

A THEORETICAL DERIVATION

A.1 Variance inducing terms in transient estimator

Here we present the derivation details and analysis for the terms that result in high variance in the baseline transient estimator (with naive NEE). Consistent with the notation used in the paper, we maintain the same symbols for clarity. For the sake of clarity, please refer to Figure 2 in our paper. The non-recursive form of the transient estimator in this infinitely-large, surface-free homogeneous scattering medium scene is expressed as follows:

$$\tilde{I} = \sum_{k=1}^N \left(\frac{W(\|\bar{\mathbf{x}}'_k\|) f_m(\boldsymbol{\omega}_k, \boldsymbol{\omega}_{k,e}) L_e}{\exp(\sigma_e \|\mathbf{x}_k - \mathbf{x}_e\|) \|\mathbf{x}_k - \mathbf{x}_e\|^2} \times \prod_{i=1}^k \underbrace{\frac{\sigma_s \exp(-\sigma_e d_{i-1})}{p_t(d_{i-1})}}_{\text{transmittance sampling}} \prod_{j=1}^{k-1} \underbrace{\frac{f_m(\boldsymbol{\omega}_{j-1}, \boldsymbol{\omega}_j)}{p_d(\boldsymbol{\omega}_j)}}_{\text{phase function sampling}} \right) \quad (1)$$

The provided equation encompasses the two main sampling PDFs in the volumetric rendering: transmittance sampling (distance) and phase function sampling (direction). As mentioned in the main paper, the steady state path construction often results in low-quality paths and many existing works continue to build paths in the same way. Considering the baseline exponential sampling and phase function sampling in the steady state renderer, we can further simplify Equation (1) as follows:

$$\tilde{I} = \sum_{k=1}^N \frac{W(\|\bar{\mathbf{x}}'_k\|) f_m(\boldsymbol{\omega}_{k-1}, \boldsymbol{\omega}_{k,e}) L_e}{\exp(\sigma_e \|\mathbf{x}_k - \mathbf{x}_e\|) \|\mathbf{x}_k - \mathbf{x}_e\|^2} \left(\frac{\sigma_s}{\sigma_e} \right)^k \quad (2)$$

$$p_t(d) = \sigma_e \exp(-\sigma_e d) \quad (3)$$

$$p_d(\boldsymbol{\omega}) = f_m(\boldsymbol{\omega}, \boldsymbol{\omega}') \quad (4)$$

Where $\boldsymbol{\omega}'$ is the ray direction before scattering. Sampling distributions employed in local importance sampling cancel some of the terms in Equation (1) out. Assuming the scattering albedo $\alpha = \sigma_s/\sigma_e$ close to 1, which holds for scattering-dominant media, we are left with the following terms for each full path:

$$\frac{W(\|\bar{\mathbf{x}}'_k\|) f_m(\boldsymbol{\omega}_{k-1}, \boldsymbol{\omega}_{k,e})}{\exp(\sigma_e \|\mathbf{x}_k - \mathbf{x}_e\|) \|\mathbf{x}_k - \mathbf{x}_e\|^2} \quad (5)$$

These terms are not perfectly cancelled out by the sampling PDFs, thereby introducing variance into the system. It is evident that these terms originate from shadow connection, encompassing connection transmittance,

directional radiance changes, inverse square falloff for emission power and the path length constraints. The transmittance and phase function terms already introduce variance, owing to the complex joint distribution of \mathbf{x}_k and $\boldsymbol{\omega}_{k,e}$. Additionally, the inverse square term is unbounded. This unbounded function maps the random variable so that the mean and standard deviation exhibit poor convergence and may potentially diverge. In the end, $W(\|\bar{\mathbf{x}}\|)$ causes the rejection of path samples, further exaggerating the variance problem.

Figure I depicts the distribution of the mean values and the stand deviation of the mean values for the inverse square term. In this 2D numerical experiment, we start to construct random paths from point $(0, 0)$, with the point emitter located at $(1, 1)$. The initial direction for all the paths is $(1, 0)$. Paths are recursively constructed, similar to the process of 3D space path tracing. The inverse square term for 500k different paths is calculated to obtain a mean value. It can be seen that single scattering (involving a single distance sample) has stable mean values. While for multiple scattering (more than 1 bounce), the mean value fails to converge, and the standard deviation of the mean value steadily increases.

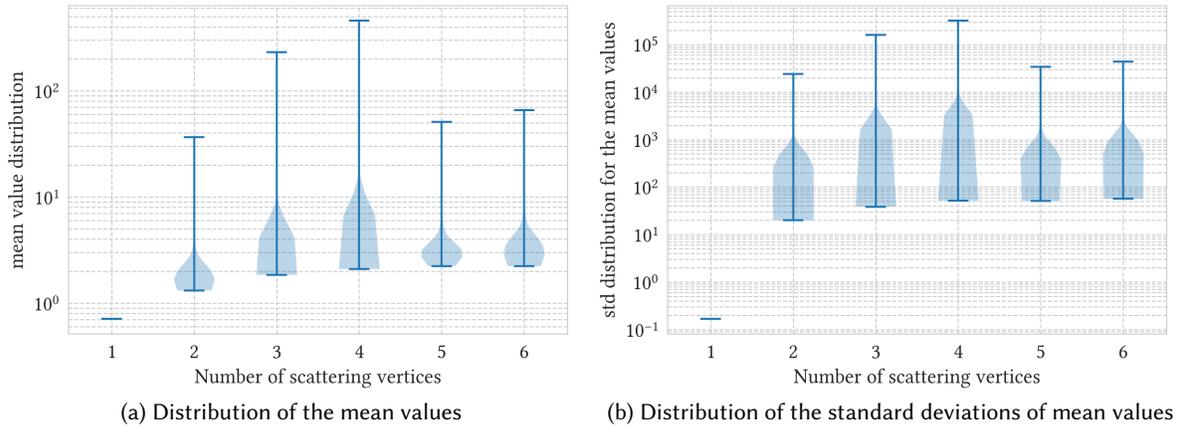


Fig. 1. Mean value distribution of the inverse square term (log scale). The mean value of the single scattering event converges, while higher bounces results in unstable mean values (a). The standard deviation for the mean value is even more unstable, meaning that the mean values have poor convergence and the original paths can have extremely high variance.

A.2 Elliptical Sampling Derivation

The complete derivation of elliptical sampling is presented in this section, accompanied by an illustration in Figure II. The sampling process initiates at \mathbf{x}_k . As outlined in the paper, the residual time range is represented by $[T_{res,m}, T_{res,M}]$, forming a 3D ellipsoidal shell for sampling. Given the assumption of an isotropic scattering medium, the azimuth angle is uniformly sampled from $[-\pi, \pi)$, disregarding the visibility of the sampled direction and resulting in a 2D elliptical ring defined by the lower and upper bounds of the residual time. The zenith angle between the previous ray direction and the elliptical connection direction is determined through straightforward phase function sampling. Consequently, the vertex sampling problem boils down to sampling the polar distance to ensure the control vertex falls within the elliptical ring. As stated in the paper, the valid sampling area (elliptical ring) in Figure II is defined by $[cT_{res,m}/\eta, cT_{res,M}/\eta)$. Therefore, we uniformly sample S from $[cT_{res,m}/\eta, cT_{res,M}/\eta)$ to obtain t . Referring to Figure IIa, t can be obtained as follows:

$$\begin{aligned}
(t - C \cos \theta)^2 + C^2 \sin^2 \theta &= (S - t)^2 \rightarrow \\
t^2 + C^2 - 2Ct \cos \theta &= S^2 - 2St + t^2 \rightarrow \\
2St - 2Ct \cos \theta &= S^2 - C^2 \rightarrow \\
t &= \frac{S^2 - C^2}{2S - 2C \cos \theta}
\end{aligned} \tag{6}$$

In a similar way, according to Figure IIb, we have:

$$\begin{aligned}
(C - t \cos \theta)^2 + t^2 \sin^2 \theta &= (S - t)^2 \rightarrow \\
t &= \frac{S^2 - C^2}{2S - 2C \cos \theta}
\end{aligned} \tag{7}$$

However, S is distributed according to $U[cT_{\text{res},m}/\eta, cT_{\text{res},M}/\eta)$, yet we actually want to sample t . Therefore, the measure conversion from S to t should be performed to obtain the correct PDF for t , since we are not integrating over the space of S but t . That is:

$$\begin{aligned}
\frac{dP}{d\mu_S}(S) &= \frac{dP}{d\mu_t}(t) \frac{d\mu_t(t)}{d\mu_S(S)} \rightarrow \\
p(t) &= p(S) \frac{d\mu_S(S)}{d\mu_t(t)}
\end{aligned} \tag{8}$$

The Jacobian term can be easily obtained through the relationship between S and t given by Equation (6). Writing S as a function of t , we have:

$$\begin{aligned}
2tS - 2Ct \cos \theta &= S^2 - C^2 \xrightarrow{\text{two roots for } S} \\
S &= \begin{cases} t - \sqrt{t^2 - 2Ct \cos \theta + C^2}, & \text{discarded since } S < 0 \text{ is not allowed} \\ t + \sqrt{t^2 - 2Ct \cos \theta + C^2}, & \text{valid solution} \end{cases}
\end{aligned} \tag{9}$$

Taking the derivative of t , we have:

$$\frac{d\mu_S(S)}{d\mu_t(t)} = \frac{dS}{dt} = \frac{t - C \cos \theta}{\sqrt{t^2 - 2Ct \cos \theta + C^2}} + 1 \tag{10}$$

According to the law of cosines, the above equation can be simplified to further reduce computation. Specifically, the denominator of the fractional part can be simplified as:

$$\begin{aligned}
\sqrt{t^2 - 2Ct \cos \theta + C^2} &= \sqrt{(S - t)^2} = S - t \rightarrow \\
\frac{d\mu_S(S)}{d\mu_t(t)} &= \frac{dS}{dt} = \frac{S - C \cos \theta}{S - t}
\end{aligned} \tag{11}$$

Therefore we have:

$$p(t) = p(S) \left(\frac{dt}{dS} \right)^{-1} = \frac{S - t}{(S - C \cos \theta)} p(S), \tag{12}$$

In the paper, $p(S)$ is given by a truncated exponential distribution, please refer to Equation (22) in the paper. This conversion is presented in its most straightforward form. Given that elliptical sampling is applied per path vertex (including the sensor vertex), it is imperative to maximize the reuse of computed values to minimize additional overhead.

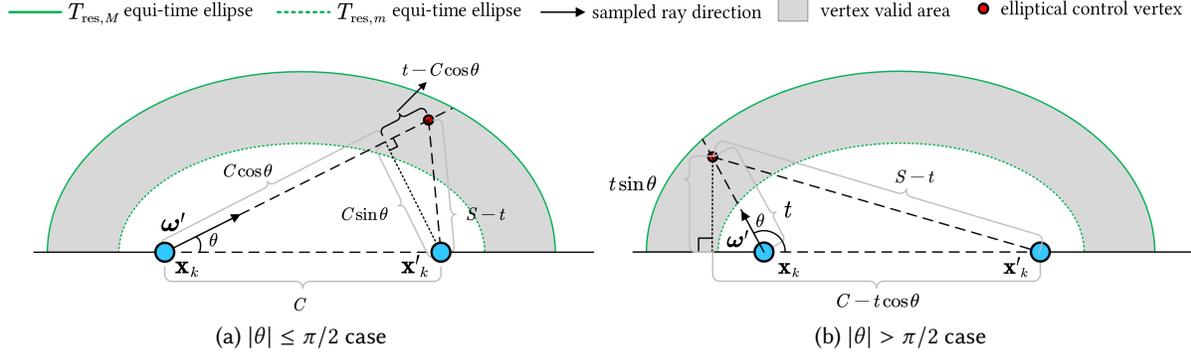


Fig. II. Two cases for elliptical sampling. The derivations for different θ s (whether the absolute value exceeds $\pi/2$) are different yet lead to the same result.

A.3 Optimality of path length control method

Here we present a straightforward proof for the necessity of path length control method. Consider a more general case, where the full path \bar{x} is sampled from a given path sampling distribution over the sampling space \mathcal{P} and $p(\bar{x})$ is its corresponding PDF. We denote $F(\bar{x})$ as the original path contribution function. When ignoring the visibility term, we assume that the support set of $F(\bar{x})$ spans the entire path space Ω . Further, when the path length constraints are applied by temporal response function $W(\|\bar{x}\|)$, the time-resolved path contribution function is denoted by $F'(\bar{x})$, and the support set is therefore denoted by $\Omega' \subset \Omega$, i.e:

$$F'(\bar{x}) = W(\|\bar{x}\|)F(\bar{x}) = 0, \text{ if } \bar{x} \in \Omega \setminus \Omega' \quad (13)$$

Note that we impose the following prerequisites on the problem: (1) all the sets and subsets in the later discussion have non-zero measure (for example, we exclude sets with only isolated points), because any non-empty set with zero measure would have zero probability of being sampled. (2) The same measure is used for different sets and no measure conversion is considered.

Then, the following proposition should hold: the necessary condition for Monte Carlo estimation of the time resolved path contribution function $F'(\bar{x})$ to be **both optimal and unbiased** is that the support set \mathcal{P}_s of the sampling distribution PDF $p(\bar{x})$ equals the support set Ω' of $F'(\bar{x})$. **Proof:**

A.3.1 Unbiasedness. We first prove that, for the estimator to be unbiased, $\Omega' \subset \mathcal{P}_s$ must hold, that is, for any given non-empty subset Q , $Q \subset \Omega'$, $Q \subset \mathcal{P}_s$. Assuming that there exists a subset denoted by Q' , $Q' \subset \Omega'$, and $Q' \subset \mathcal{P} \setminus \mathcal{P}_s$, meaning that any path $\bar{x}' \in Q'$, $p(\bar{x}')$ is zero. So the actual integral result I is given by:

$$I = \int_{Q'} F(\bar{x}) d\mu(\bar{x}) + \int_{\Omega' \setminus Q'} F(\bar{x}) d\mu(\bar{x}) \quad (14)$$

where $\int_{Q'} F(\bar{x}) d\mu(\bar{x}) \neq 0$

and the expectation $\hat{\mu}$ of the Monte Carlo integration estimate is given by:

$$\hat{I} = \int_{\mathcal{P}_s} \frac{F(\bar{x})}{p(\bar{x})} p(\bar{x}) d\mu(\bar{x}) = \int_{\Omega' \setminus Q'} F(\bar{x}) d\mu(\bar{x}) \neq I \quad (15)$$

which indicates that the estimator is biased. This results in a conflict. Therefore, when the estimator is unbiased, $\Omega' \subset \mathcal{P}_s$ must hold.

A.3.2 Optimality. We then prove that, for the estimator to be optimal, $\mathcal{P}_s \subset \Omega'$ should hold. We begin by constructing the following sampling distribution with PDF denoted as $p^*(\bar{\mathbf{x}})$. This sampling distribution forms an unbiased estimate of I , and the support set \mathcal{P}_s^* of the PDF has the following properties: $\Omega' \subset \mathcal{P}_s^*$ (unbiasedness), $\mathcal{P}_s^* \not\subset \Omega'$. Now we prove that $p^*(\bar{\mathbf{x}})$ can not be optimal since through this distribution, we can construct a better estimator with lower variance. The variance of the \hat{I} is given by:

$$\text{Var}(\hat{I}) = \int_{\Omega'} \frac{f^2(\bar{\mathbf{x}})}{(p^*(\bar{\mathbf{x}}))^2} p^*(\bar{\mathbf{x}}) d\mu(\bar{\mathbf{x}}) - I^2 = \int_{\Omega'} \frac{f^2(\bar{\mathbf{x}})}{p^*(\bar{\mathbf{x}})} d\mu(\bar{\mathbf{x}}) - I^2 \quad (16)$$

We denote $Q = \mathcal{P}_s^* \setminus \Omega'$. Note that $p^*(\bar{\mathbf{x}})$ is a valid PDF, we have:

$$\int_{\mathcal{P}_s^*} p^*(\bar{\mathbf{x}}) d\mu(\bar{\mathbf{x}}) = 1 \quad (17)$$

$$\int_Q p^*(\bar{\mathbf{x}}) d\mu(\bar{\mathbf{x}}) = \beta \in (0, 1) \quad (18)$$

So, we can construct a sampling distribution with PDF $q^*(\bar{\mathbf{x}})$ as the following form:

$$q^*(\bar{\mathbf{x}}) = \begin{cases} \frac{1}{1-\beta} p^*(\bar{\mathbf{x}}), & \bar{\mathbf{x}} \in \Omega' \\ 0, & \text{otherwise} \end{cases} \quad (19)$$

Note that this is still a valid PDF: it is easy to verify that $q^*(\bar{\mathbf{x}})$ integrates to 1 in its domain, since $p^*(\bar{\mathbf{x}})$ integrates to $1 - \beta$ on set Ω' . With this sampling distribution (and its PDF), the variance now becomes:

$$\text{Var}'(\hat{I}) = \int_{\Omega'} \frac{(1-\beta)f^2(\bar{\mathbf{x}})}{p^*(\bar{\mathbf{x}})} d\mu(\bar{\mathbf{x}}) - I^2 < \text{Var}(\hat{I}) \quad (20)$$

which is lower than the variance of the original estimator with PDF $p^*(\bar{\mathbf{x}})$. Therefore, we proved that for an unbiased estimator to be optimal, $\mathcal{P}_s \subset \Omega'$ must hold. Then we can prove that $\mathcal{P}_s = \Omega'$. The necessary condition of unbiasedness and optimality is therefore proved.

This proposition shows that: to achieve the optimal unbiased estimate of the time-of-flight radiance, the support set of the sampling path PDF must align with the support set of the path contribution function, in other words, it is essential to avoid getting samples with zero contribution. Normally, without path length control method, the estimator retains its unbiasedness (if it is unbiased), yet only $\Omega' \subset \mathcal{P}_s$ holds. Intuitively, Ω' may occupy only a fraction of the total volume of \mathcal{P}_s , implying that, for a given sample $\bar{\mathbf{x}}$, it is highly probable to belong to $\mathcal{P}_s \setminus \Omega'$ and has zero contribution. Therefore, this results in massive sample rejection as the majority of samples yield zero contribution.

Note that this proposition does not guarantee the sufficient condition. Being *unbiased* imposes restrictions on the use of non-delta blur in both spatial and temporal domain, and being *optimal* requires that the sampling PDF accurately approximates the shape of $F'(\bar{\mathbf{x}})$ to maintain a constant ratio between them. This can not be achieved without zero-variance random walk theory and data-driven methods. Path control method, however, does guarantee the necessary condition for optimality by aligning the support set of the integrand with the sampling PDF for Monte Carlo integration. Therefore, it establishes the necessary condition for optimality and avoids sample rejection.

B IMPLEMENTATION DETAILS

In this section, the details about our implementation are given, encompassing: (1) The modifications made in pbrt-v3 [Pharr et al. 2023] and Tungsten [Bitterli 2018; Liu et al. 2022]. (2) The implementation differences between path tracing method and photon based methods.

B.1 Modifications of Existing Renderers

The implementation of the proposed methods is based on the open-source renderers pbrt-v3 [Pharr et al. 2023] and Tungsten [Bitterli 2018]. We primarily utilized pbrt-v3 for its unidirectional volumetric path tracing (UDPT), and for Tungsten, we employed the modified version introduced by Liu et al. [2022], where they implemented their temporally sliced photon primitive methods. Their modified version enables the rendering of scenes under time-of-flight settings to some extent. However, these renderers are initially designed for camera warped, full transport rendering settings, thereby additional modifications are necessary to align them with the requirements of our paper. The major modifications are outlined below:

B.1.1 Camera warped settings. We noted that pbrt-v3 [Pharr et al. 2023] does not inherently support rendering time-resolved radiance. We modified its volumetric path tracing and bidirectional path tracing (BDPT) modules to support camera-warped. Regarding the Tungsten renderer, Liu et al. [2022] did implement the time-of-flight rendering method in the original steady state rendering Tungsten [Bitterli 2018]. However, their temporally sliced photon primitives do not retain primitive shapes in camera warped settings. Additionally, the scene-to-sensor transport time is overlooked in their implementation. To rectify this, we developed code to calculate the scene-to-sensor transport time, involving ray-primitive intersection and distance calculation. Due to complexity of the photon primitives, the modification in the code are non-trivial (such as beams and planes). Note that certain estimators, such as photon balls, volumes and photon hyper-volumes are abandoned in our implementation due to incorrect Jacobians needed in sampling PDF for camera-warped settings, and the further mathematical correction to address this problem is beyond the scope of this paper.

B.1.2 Transient rendering. The transient rendering in Liu et al. [2022] is based on frame-by-frame time-gated rendering. Therefore, expensive operations like parameter parsing, scene loading and BVH construction are performed in each frame without any reuse. Since most existing transient rendering approaches rely on the temporal path reuse, we implemented a transient path reuse framework based on the code of Liu et al. [2022]. This modified transient rendering supports path tracing, photon points and photon beams methods. We also implemented the same framework in pbrt-v3 to support its UDPT and BDPT.

B.1.3 Full transport simulation. pbrt-v3 inherently supports full transport simulation, while the work of Liu et al. [2022] focuses on volumetric transport, leaving the surface transport part not directly usable. To address this, we modified the code to enable time-resolved surface transport simulation. However, the modification is achieved through the time-resolved adaptation for the photon mapping of the original renderer, meaning that there exists a gap between the surface transport and the volumetric transport. It's worth pointing out that the steady state higher-dimensional photon primitives [Bitterli and Jarosz 2017; Deng et al. 2019] face a similar issue as they do not focus on surface transport.

B.1.4 Guided homogeneous medium. The class GuidedHomogeneousMedium is implemented in the medium modules of both pbrt-v3 and Tungsten. Note that Tungsten supports non-exponential transmittance, but our implementation is focused on the most prevalent exponential transmittance model. This medium class includes member methods for elliptical sampling and DA-based distance sampling. Notably, this medium implementation degrades to the standard homogeneous medium implementation when both DA-based distance sampling and elliptical sampling are disabled. Hence, throughout our paper, we used this class instead of the baseline implementations. Both pbrt-v3 and Tungsten share two common hyper-parameters, which control whether to enable DA-based distance sampling and elliptical sampling, respectively. However, in Tungsten, we introduce an additional hyper-parameter named *strict-time-mode*, which will be further discussed in Section B.2. We optimize our implementation by incorporating Intel's SSE instructions to accelerate RIS, necessitating an upgraded of the entire old pbrt-v3 to C++17 to support this feature.

B.1.5 Integrated DA phase function. The class `IntegratedDAPhaseFunction` is implemented in the phase function modules of both `pbirt-v3` and `Tungsten`. This phase function module will degrade to `Henry-Greenstein` phase function class if either the 3D table or the elliptical sampling struct is not provided to this class. If correctly initialized with precomputed 3D table, this class will provide EDA direction sampling with MIS supported.

B.1.6 Matching the rendering outputs. We made adjustments to certain material implementations in `Tungsten` to align its rendering outputs with those of `pbirt-v3`. The most substantial modifications are applied to the coated material and microfacet surface models, given their frequent use in our experiments. The plastic material (class `PlasticMaterial`) in `pbirt-v3` should work similarly to the rough-plastic material (class `RoughPlasticBsdf`) in `Tungsten`. Note that there is a `PlasticBsdf` class in `Tungsten`, yet it does not account for surface roughness (unlike the plastic material in `pbirt-v3`). The main difference in this material lies in how the Fresnel effect is treated: `pbirt-v3` only accounts for the Fresnel effect for the specular part, while `Tungsten` weights both the specular and diffuse part by the Fresnel coefficient. The implementation of GGX microfacet distribution model also differs in these two renderers, so we implement GGX model of the `pbirt-v3` in `Tungsten`. In all, the modifications encompass the microfacet model, rough plastic surface model, rough dielectric BSDF and the default UV-coordinate calculation procedures for meshes without UV-coordinates. These modifications ensure consistent rendering results between the two different renderers.

B.2 Differences between Path Tracing and Photon Based Methods

Our method is initially derived in path tracing frameworks, therefore it can be directly applied to path tracing based methods. In our paper, we have also demonstrated the versatility of our method that it can be applied to the photon based methods. Since the photon based methods can be considered as bidirectional methods, their implementation can be quite different. Here we present the pseudo code of the complete transport sampling process and highlight the difference between these two frameworks.

The pseudo code presented in Algorithm 1 describes the time range subdivision and target time sampling procedures, before actual path construction and connection. Here we use two different notation for the subdivided object: when the program is used for transient rendering then the subdivided objects are referred to as *frames*; when for rendering time-gated images where sampling according to the transient response weight $W(\|\bar{x}\|)$ is needed, the time range is first subdivided (please refer to Section 5.2 of our paper), and in this situation, the objects are denoted as *bins*. This time-of-flight path tracing procedure unifies transient rendering and time-gated rendering: for transient rendering, the sample function in line 6 is a uniform or low discrepancy sampler to distribute the full path samples uniformly in time domain, while for time-gated rendering, the sample function can be manually implemented to best capture the shape of $W(\|\bar{x}\|)$. It is worth noting that the existing works typically ignore the subdivision and target time sample part.

Algorithm 2 provides the pseudo code of time-of-flight volumetric path tracing with our DARTS path sampling method. There are two points to be noted:

On line 12, we directly use ω_j, d_m^j to perform elliptical sampling, instead of sampling a new ray from the current position. As mentioned in Section 5.1 of our paper, this reuse has no impact on rendering quality, yet for each elliptical-sampled connection, no extra ray-intersection is required. Therefore, this reuse lifts the extra ray-intersection of elliptical sampling, which can be computationally burdensome in scenes with complex geometries.

From line 20 to line 30, it can be seen that the surface pass still retains direct shadow connection, since we are unable to deterministically put a surface control vertex, while the medium shadow connection is replaced by our elliptical sampling. Also, elliptical sampling does not restrict the connected vertices to be medium vertices.

Algorithm 1 Time-of-flight Path Tracing

```

1: procedure TOFPATHTRACING( $T_{t,m}, T_{t,M}, \mathbf{x}_0, \boldsymbol{\omega}_1, \mathcal{S}, N, k, n_f$ )
2:    $\triangleright \mathbf{x}_0$  ... Camera position, the starting point of a ray
    $\boldsymbol{\omega}_1$  ... Initial ray direction sampled from the camera model
    $[T_{t,m}, T_{t,M})$  ... Time range to be rendered
    $\mathcal{S}$  ... Scene settings: such as camera, participating media information and geometries, etc.
    $N, k, n_f$  ... Number of path samples, max bounces and number of frames (bins), respectively
3:    $\mathcal{L} \leftarrow \{L_1, \dots, L_{n_f}\} = \text{Subdivision}(T_{t,m}, T_{t,M}, n_f)$   $\triangleright$  subdividing the time range into  $n_f$  frames
4:    $\mathcal{B} \leftarrow \{w_1, \dots, w_{n_f}\} = W(\mathcal{L})$   $\triangleright$  Evaluating  $W(\|\bar{\mathbf{x}}\|)$  to get sampling weights
5:   for  $n = 1, \dots, N$  do
6:      $[T_{t,m}^i, T_{t,M}^i, p_i] \leftarrow \text{Sample}(\mathcal{L}, \mathcal{B})$   $\triangleright$  Sample a time bin according to bin weight with probability  $p_i$ 
7:      $L \leftarrow \text{PathIntegral}(\mathcal{S}, k, T_{t,m}^i, T_{t,M}^i, \mathbf{x}_0, \boldsymbol{\omega}_1)$   $\triangleright$  Get full path samples in time range  $[T_{t,m}^i, T_{t,M}^i)$ 
8:      $L_i \leftarrow L_i + L/p_i$   $\triangleright$  Accumulate path contribution (radiance) to the frame (bin)  $i$ 
9:   return  $\mathcal{L}$ 

```

In photon based methods, the sensor pass normally does not recursively trace the ray (unless the BSDF is purely forward or specular). Therefore, DARTS can only be applied to photon pass. Since the photon pass simulates the light path from the emitter, there are several differences in its implementation:

- The target vertex to be connected to is the sensor vertex $\mathbf{x}_{\text{sensor}}$ instead of \mathbf{x}_e . So, instead of sampling a point on a random emitter, we sample a point on the sensor. For pinhole perspective camera, the point can be obtained deterministically since the sensor has no physically area, while for realistic lens based cameras, points on the lens should be sampled.
- DA-based distance sampling considers the camera as an equivalent emitter and use it instead of the real emitter to calculate DA. Therefore, the DA actually solves for the flux of importons rather than photons. It can be concluded that our DA-based distance sampling can be applied as an approximation to the adjoint transport information.
- Shadow connection is not performed in photon pass. Therefore, instead of evaluating full path contribution after performing elliptical sampling, the sampled control vertices are stored into the photon map. The path vertex is stored in the photon map only when it's a surface vertex, distinguishing it from the baseline photon tracer that records all vertices. This method ensures that all photon records in the photon map have a much higher probability, as photon mapping involves gathering, to form full paths that satisfy the path length constraints.
- The length of the full path is available only when the photon pass and the sensor pass are completed. We only record the power, time-of-flight and position, etc. in a photon record. The calculation of the full path time slightly differs: the photon points are gathered using a point-beam 2D blur estimator [Křivánek et al. 2014], therefore, the scene-to-sensor transport time is given by the projected length of a photon point, instead of the Euclidean distance.
- The *strict-time-mode* hyper-parameter mentioned in Section B.1.5 becomes relevant when dealing with scenes containing pure specular materials, in this case, determining the potential path time of a photon by connecting it directly to the sensor becomes impractical. The specular surfaces introduce alternative paths that need at least one more specular bounce. Therefore, discarding all photons that can not satisfy the path length constraints when connected to the sensor may bias the rendering. We provide the option to set *strict-time-mode* to false, allowing users to disable the photon selection. When set to false, photons with a full path time $T < T_{t,m}^i$ will not be discarded, therefore preserving the specular effect.

Algorithm 2 Time resolved path integral (volumetric path tracing)

```

1: procedure PATHINTEGRAL( $T_{t,m}^i, T_{t,M}^i, \mathbf{x}_0, \boldsymbol{\omega}_1, \mathcal{S}, k$ )
2:    $\triangleright [T_{t,m}^i, T_{t,M}^i)$  ... Target time range sampled, all full path samples should fall into this range.  $\triangleleft$ 
3:    $[\mathbf{x}_e, p_e] \leftarrow \text{SampleEmitter}(\mathcal{S})$   $\triangleright$  Sample a emitter vertex on the emitter with probability  $p_e$ 
4:    $[\mathbf{x}_s^1, d_m^1] \leftarrow \text{RayIntersect}(\mathcal{S}, \mathbf{x}_0, \boldsymbol{\omega}_1)$ 
5:    $\beta \leftarrow 1, L \leftarrow 0$   $\triangleright$  Path throughput and accumulated contribution initialization
6:    $T_e \leftarrow 0$   $\triangleright$  Path elapsed time initialization
7:   for  $j = 1, \dots, k$  do
8:      $T_{\text{res}} \leftarrow \text{ResidualTimeSampling}(T_{t,m}^i, T_{t,M}^i, T_e)$ 
9:     if  $T_{\text{res}} \leq 0$  then
10:       $\triangleright$  break  $\triangleright T_e + T_{\text{res}} > T_{t,M}^i$ : the subsequent bounces will break the path length constraints
11:       $\triangleright$  Elliptical sampling: Equation (22), (23) in the paper
12:      note that the ray direction and intersection are reused, see section 5.1 in the paper  $\triangleleft$ 
13:       $[\mathbf{x}_{\text{ell}}, \beta_{\text{ell}}] \leftarrow \text{EllipticalSampling}(\mathcal{S}, \mathbf{x}_{j-1}, \boldsymbol{\omega}_j, T_{\text{res}}, d_m^j)$ 
14:      if  $\beta_{\text{ell}} > 0$  then
15:         $L \leftarrow L + \beta \times \beta_{\text{ell}} \times \text{ShadowConnection}(\mathbf{x}_e, \mathbf{x}_{\text{ell}}, \mathcal{S})$ 
16:         $\triangleright$  DA distance sampling: Equation (17) - (20) in the paper  $\triangleleft$ 
17:         $[t, \Delta\beta, \text{is\_surface}] \leftarrow \text{DistanceSampleDA}(\mathcal{S}, \mathbf{x}_{j-1}, \boldsymbol{\omega}_j, T_{\text{res}}, d_m^j)$ 
18:         $T_e \leftarrow T_e + \eta t / c$   $\triangleright$  Update path elapsed time
19:         $\beta \leftarrow \beta \times \Delta\beta$ 
20:         $\mathbf{x}_j \leftarrow \mathbf{x}_{j-1} + \boldsymbol{\omega}_j t$ 
21:        if  $\text{is\_surface} = \text{true}$  then
22:           $\triangleright$  Surface direction sampling: direct shadow connection is performed
23:          then we sample scattered ray direction  $\triangleleft$ 
24:           $L_s \leftarrow \text{ShadowConnection}(\mathbf{x}_e, \mathbf{x}_s^j, \mathcal{S}) \times \text{EvalCosineTerm}(\mathcal{S}, \mathbf{x}_s^j)$ 
25:          if  $T_e + \eta \|\mathbf{x}_s^j - \mathbf{x}_e\| / c \in [T_{t,m}^i, T_{t,M}^i)$  then
26:             $L \leftarrow L + \beta L_s / p_e$ 
27:             $[\boldsymbol{\omega}_{j+1}, \beta_s] = \text{BSDFSampling}(\mathcal{S}, \mathbf{x}_s^j, \boldsymbol{\omega}_j, \mathbf{x}_e)$ 
28:             $\beta \leftarrow \beta \times \beta_s$ 
29:          else
30:             $\triangleright$  Medium direction sampling: direct shadow connection is disabled  $\triangleleft$ 
31:             $\triangleright$  In EDA sampling, one-sample-model MIS will randomly sample one approach to be used from
32:            (EDA direction sampling) and (phase function sampling) and combine both approaches with
33:            balance heuristic  $\triangleleft$ 
34:             $[\boldsymbol{\omega}_{j+1}, \beta_p] = \text{EDASampling}(\mathcal{S}, \mathbf{x}_j, \boldsymbol{\omega}_j, T_{\text{res}}, \mathbf{x}_e)$ 
35:             $\beta \leftarrow \beta \times \beta_p$ 
36:          if  $\beta = 0$  then
37:             $\triangleright$  Break if the throughput is 0
38:             $\triangleright$  break
39:             $[\mathbf{x}_s^{j+1}, d_m^{j+1}] \leftarrow \text{RayIntersect}(\mathcal{S}, \mathbf{x}_j, \boldsymbol{\omega}_{j+1})$ 
40:          return  $L$ 

```

For photon primitive estimators in higher dimensions, such as beams, our elliptical sampling can not be directly applied in rendering, since those estimators do not gather individual photons but interpolate between consecutive photons (on a path) via ray-primitive intersection. In those estimators, elliptical sampling is only

used to determine whether a photon path is longer than needed. DA-based distance sampling, however, can still be used to improve the path construction quality, and we will be presenting some results in the following section.

B.3 Tabulated Sampling

The EDA direction sampling proposed in Section 5.1 of our paper employs offline tabulation for inverse transform sampling. As stated in the paper, tabulation generates a 3D table. Throughout the experiments, we use (128, 128, 256) as table shape, meaning that the dimensions for C/S and S are uniformly subdivided into 128 bins, while the last dimension for $\cos \theta$ is binned by 256. We refer to the first two dimensions combined as *grid* for clarity. Firstly, we will uniformly sample the grid to obtain some C and S samples for different ellipses. Then we evaluate the elliptical polar distance for each $\cos \theta$ samples generated by stratified sampling from the 256 bins. Then, for each bin in the $\cos \theta$ dimension, we approximate the integral given in Equation (16) of the paper via Monte Carlo integration, with the corresponding elliptical polar distance as the upper bound of integration. For each grid, a total of 4096 different ($C/S, S, \cos \theta$) samples are drawn, and it is already enough to produce smooth tabulation results. After MC integration, we take the prefix sum of the $\cos \theta$ dimension and normalize the result so that the last bin equals to 1 (valid cumulative distribution function).

Note that the sampling and evaluation can be naturally structured in a single-instruction-multiple-thread (SIMT) way, we therefore choose to implement this tabulation on GPU. The implementation is based on PyTorch 2 [2023], and the parallel computation (one row in the grid at a time) is achieved via the GPU tensor computation. The code is compiled to achieve further acceleration, resulting in a 5 second computation time on a single entry-level GPU (Nvidia TITAN RTX). For more detail on the implementation, one can refer to `da_ell/tabulation.py` in our provided code of modified *pbrt-v3*.

B.4 Reproduction of the Uniform-Time Sampling Method

As stated in Section 6.1 of the paper, we reproduce and compare the work of Jarabo et al. [2014]. Since their method was originally designed for bidirectional path tracing and transient rendering, several modifications are made to adapt it for comparison with unidirectional path tracing and time-gated rendering rendering setup.

- We implement only the sampling method described in Section 5 of their paper, specifically: exponential sampling, line-to-point sampling and time angular sampling. Methods related to temporal kernel density estimation and photon mapping method are not included.
- The original time-angular sampling method (Section 5.3 of their paper) appears to be flawed. Equation S.98 in their supplementary note has a missing r_{i+1} in the numerator. Additionally, their sampling formula considers only the case where $|\mathbf{l}| \geq r_{i+1}$, whereas $|\mathbf{l}| < r_{i+1}$ is also possible. Therefore, their official implementation (2019 latest version) is incorrect. We have therefore developed a corrected version, which we briefly introduce the correct derivation here.

Equation S.98 in their supplementary note should be corrected to the following form. For simplicity, we omit the η/c term by setting the speed of light and the relative refraction index to 1:

$$\theta(t) = \arccos \left(\frac{|\mathbf{l}|^2 + 2tr_{i+1} - t^2}{2|\mathbf{l}|r_{i+1}} \right) \quad (21)$$

As for t , we should consider two cases:

$$t(\xi) = \begin{cases} (2r_{i+1} + |\mathbf{l}| - 2r_{i+1} + |\mathbf{l}|)\xi + 2r_{i+1} - |\mathbf{l}|, & |\mathbf{l}| < r_{i+1} \\ (2r_{i+1} + |\mathbf{l}| - |\mathbf{l}|)\xi + |\mathbf{l}|, & \text{otherwise} \end{cases} \quad (22)$$

These two cases result in different sampling formulas and thus require an if-branch to be correctly incorporated (details can be found in our open-source code). Jarabo et al. [2014] only considered the second case of the above

equation. Both cases can be derived using a simple triangle plot. While Jarabo et al.'s derivation contains flaws, their theory remains effective and inspiring.

C EXTRA RESULTS

C.1 Ablation Study

In this section, we demonstrate the efficacy of the proposed method through an ablation study. Within the CORNELL BOX scene, 4 distinct rendering settings are showcased: baseline (without any modification), elliptical sampling (enabling elliptical sampling only), DA distance sampling (enabling DA distance sampling only), and full DARTS (both enabled). Results are presented in Figure III. In the figures, the scattering coefficients σ_s are set to 0.2m^{-1} and 0.4m^{-1} for the first two and last two rows (only in the bounding box of the meshes), respectively.

Clearly, our proposed path construction and connection strategies significantly enhance ToF rendering quality. Furthermore, it is evident that DA-based distance sampling synergizes well with elliptical sampling. It is noteworthy that DA-based distance sampling proves more effective in scenes with optically thicker scattering media, as indicated by MSE reduction in both $\sigma_s = 0.2\text{m}^{-1}$ and $\sigma_s = 0.4\text{m}^{-1}$ cases.

C.2 Rendering results in other scenes

The rendering results in two other scenes are given. The BATHROOM scene [Bitterli 2016] features mirror specular, glossy metallic surface and rough plastic substrate material (see Figure IV). The KITCHEN scene [Bitterli 2016] features glossy metallic reflection (see Figure V), mainly. In both scenes, photon based methods perform poorly on glossy surfaces, and this is the main factor that introduces error for DARTS PP.

C.3 Strict Time Mode & DARTS Photon Beams

In this section, we showcase images rendered using the photon beams method with our DARTS path sampling. Additionally, two illustrative figures depicting the rendering results with the *strict time mode* enabled are given. Figure VI demonstrates that DARTS can be successfully extended to photon beam methods. However, due to its limited effectiveness and tuning complexity, the photon beam with DARTS is not included in the main paper.

Furthermore, the impact of the *strict time mode* on the rendering of mirror specular components is evident in Figure VI(a) and Figure VI(b). In the GLOSSY DRAGON scene, featuring a mirror on the left wall, enabling the *strict time mode* results in the loss of the mirror reflection. Also, the absence of specular components in certain time range render the dragon statue dimmer (second row). Nevertheless, the base of the statue and the floor exhibit higher-quality rendering with fewer visual artifacts. This suggests that the *strict time mode* might be beneficial for improving rendering quality in diffuse-dominant scenes.

C.4 Comparison with other estimators or samplers

In this section, we showcase images rendered using various distance sampling strategies, including equiangular sampling [Kulla and Fajardo 2012] and uniform sampling. We provide a brief discussion on the advantages of the proposed DA-based distance sampling. Subsequently, we conduct a simple comparison between baseline bidirectional path tracing and our unidirectional path tracing approach integrated with DARTS. This comparison covers both quantitative and qualitative aspects under equal-time conditions. Note that, in this section, EDA direction sampling is turned off, in order to emphasize the comparison performance of other modules.

Comparison with other distance sampling methods is depicted in Figure VII. The last two sub-figures compares two different candidate sampling methods: equiangular sampling and truncated exponential sampling. It can be seen that, despite its success in steady-state and low-bounces settings, equiangular sampling is still outperformed by truncated distance sampling. Note that the first three images in the figure are not the original rendering outputs. We average several frames and normalize the result with their corresponding 0.99 quantile number. The

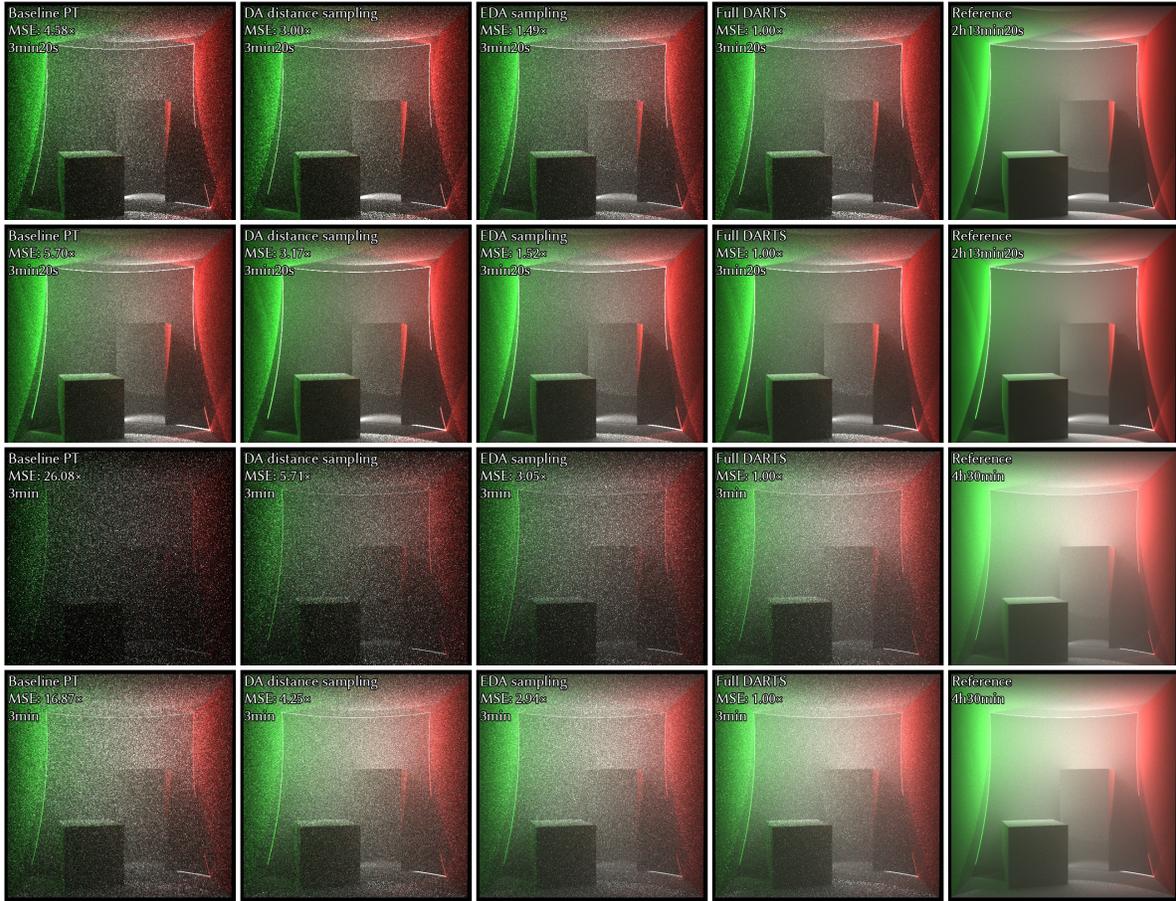


Fig. III. Ablation study: Images on the same row are rendered under the same settings. The first two rows are rendered with σ_s set to 0.2m^{-1} while the last two rows are rendered with σ_s set to 0.4m^{-1}

number of frames to be averaged varies based on the rendering time of each frame. Quantitative results (MSE) are computed in an equal-SPP manner (using original outputs to calculate MSE), while qualitative results (images) are presented in an equal-time manner, ensuring the averaged images require roughly the same time to produce. In either case, DA-based distance sampling significantly improves rendering quality.

Additionally, it is worth noting that equiangular sampling [Kulla and Fajardo 2012] can indeed enhance rendering quality, but its effectiveness is noticeable mainly in steady-state rendering tasks with a small maximum number of bounces (e.g., 2 bounces). We present the results of a simple experiment with three different scenarios: steady-state rendering (2 bounces, full transport, see Figure VIII(a)); time-gated rendering (2 bounces, full transport, see Figure VIII(b)); time-gated rendering (at most 80 bounces, full transport, see Figure VIII(c)). The first row is rendered using equiangular sampling-based path tracing (10k SPP, rendering time 440s), while the second row is rendered by DARTS PT (10k spp, rendering time 670s). Although our proposed method requires more rendering time, in multiple-scattering time-of-flight rendering scenarios, the rendering quality of DARTS PT still



Fig. IV. The camera-warped time gated rendering of the BATHROOM scene (modified version)[Bitterli 2016]. Note that DARTS PP enables *strict time mode*, therefore the mirror on the wall is dark compared to the ground truth (rendered by DARTS PT, 5.73h).



Fig. V. The camera-warped time gated rendering of the KITCHEN scene [Bitterli 2016]. For photon based methods, significant visual artifacts can be observed, since compared to path tracing, photon based methods are less competitive at rendering glossy and specular materials.

significantly exceeds that of equiangular sampling-based path tracing with doubled SPP (20k, rendering time 890s). Note that the scattering effect is relatively strong in this scene ($\sigma_s = 0.6m^{-1}$, standard CORNELL BOX scene scale).

We further compare unidirectional path tracing (UDPT) with our DARTS as a path sampling method against baseline bidirectional path tracing with path MIS (see Figure IX). Although this is not the most suitable case for BDPT, we still observe an almost three-fold improvement compared to baseline UDPT. Moreover, we achieve

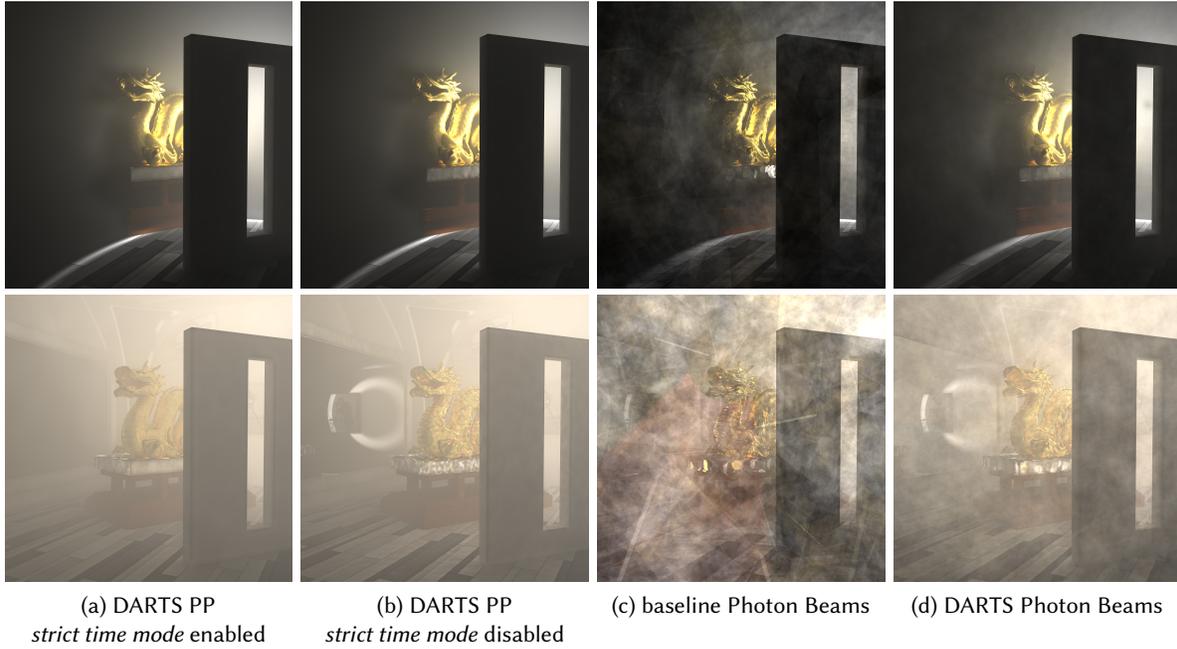


Fig. VI. Equal-time comparison (rendering time: 14m50s) in the GLOSSY DRAGON SCENE. In (a) and (b), we observe that the strict time mode brings subtle enhancements to surface scattering but completely excludes the rendering of mirror specular components. The rendering results of DARTS Photon Beams in (d) are not as effective as those of DARTS PP (refer to column (b)), hence they were not thoroughly analyzed in the primary experiments of our paper. Nevertheless, the improvement is noteworthy when compared to (c).



Fig. VII. Equal-SPP (SPP 10k) comparison in CORNELL BOX scene. The DA distance sampling is non-trivial in time-of-flight rendering tasks, and can not be easily replaced by other existing distance sampling methods.

roughly a four-fold improvement compared to BDPT. However, we also found that since BDPT performs per-vertex connection, it usually has a much lower SPP value compared to UDPT with the same rendering time. This is potentially detrimental, because in scenes with complex geometry and poor path length control, higher SPP values are needed to get more samples. In more complex scenes like GLOSSY DRAGON and STAIRCASE, where vertex connection requires ray-intersection (which is expensive in scenes with complex geometry), BDPT is found to perform even worse than baseline UDPT.

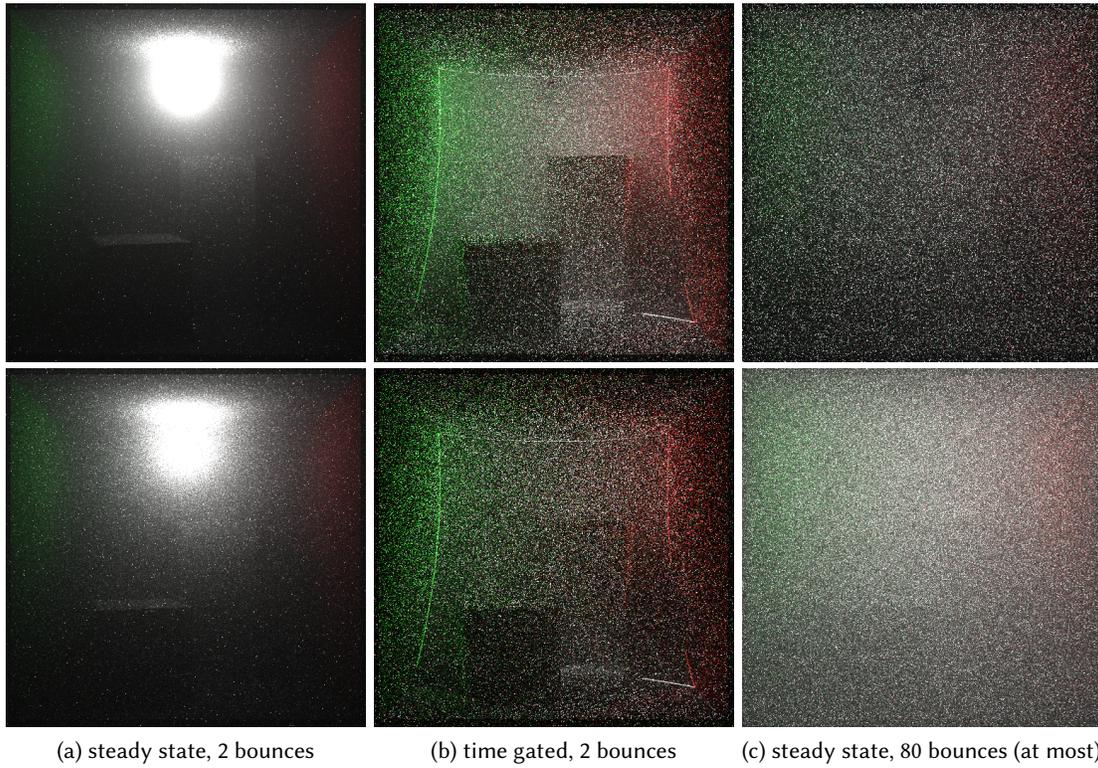


Fig. VIII. Effectiveness of equiangular sampling: equiangular sampling is not a good option in time-of-flight sensor simulation where the number of bounces can be high. The first row of the figure is the results of equiangular sampling methods, while the second row is the rendered images of DARTS PT.

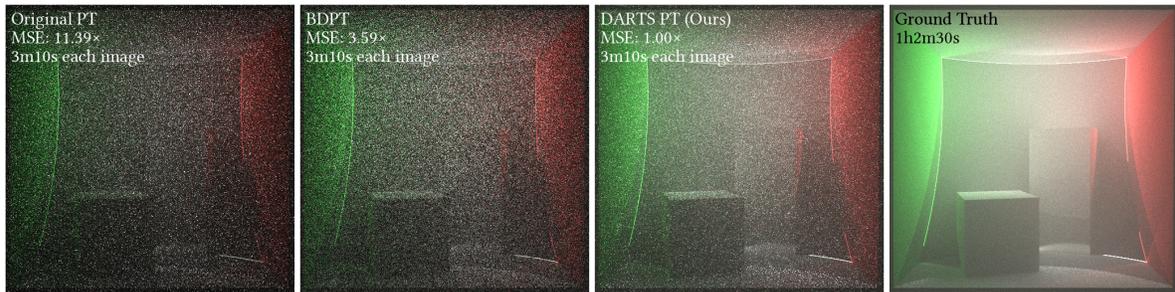


Fig. IX. Comparing unidirectional path tracing (UDPT) with bidirectional path tracing (BDPT). Under the equal-time comparison settings, bidirectional path tracing method can indeed outperform baseline UDPT, yet it is still outperformed by the UDPT equipped with DARTS.



Fig. X. Comparison with other estimators proposed by Liu et al. [2022], mainly *transient photon planes* and *transient photon volumes*. Except for the image rendered by photon volumes in the second time point of the GLOSSY DRAGON scene (row 2, col 2), of which the MSE is lower than photon points method, all other cases are found to perform poorly. Therefore, these estimators are not included in our paper.

In Figure X, we present comparison results among photon points, DARTS-based photon points, and estimators not covered in our paper (*transient photon planes* and *transient photon volumes*) proposed by Liu et al. [2022]. Photon planes estimator, despite being unbiased, introduces a heavier computational overhead due to the need to trace 'ghost bounces' and perform additional ray-primitive intersections. Furthermore, without MIS, it is susceptible to singularity problems. On the other hand, the photon volumes estimator produces incorrect images as the Jacobians derived by Liu et al. [2022] only function correctly in camera-unwarped settings. Consequently, it cannot be relied upon to generate valid rendering images.

C.5 Violation to the assumption of DA

We have also tested our method in scene settings that may challenge the assumptions of DA. Specifically, we explored two situations in a homogeneous-media setup: media with strong absorption, where we relax the assumption that $\sigma_s \gg \sigma_a$ and phase function with directionality (various g parameters for the Henyey-Greenstein phase function). For both scenarios, We compare the performance of our sampling method with that of a naive path tracer. All images are rendered with 5k SPP.

From the results in Figure XI, we can conclude that even when the low absorption assumption is violated, our algorithm can still generate results with significantly improved variance profile. This robustness is partly attributed to elliptical sampling, which will not be much affected by medium properties, especially the scattering and absorption coefficients (for phase function parameter, refer to the next section). It's obvious that the degradation rate of our methods does not escalate as rapidly as the complexity of the scene setup. However, in cases with varying directionality in Figure XII, our method can still demonstrates considerable improvements for backward scattering. In contrast, for forward scattering, there is a notable performance drop as g increases. Despite this, our method still renders images in half the time compared to other methods. We believe the performance drop is related to the MIS between EDA direction sampling and phase function sampling.

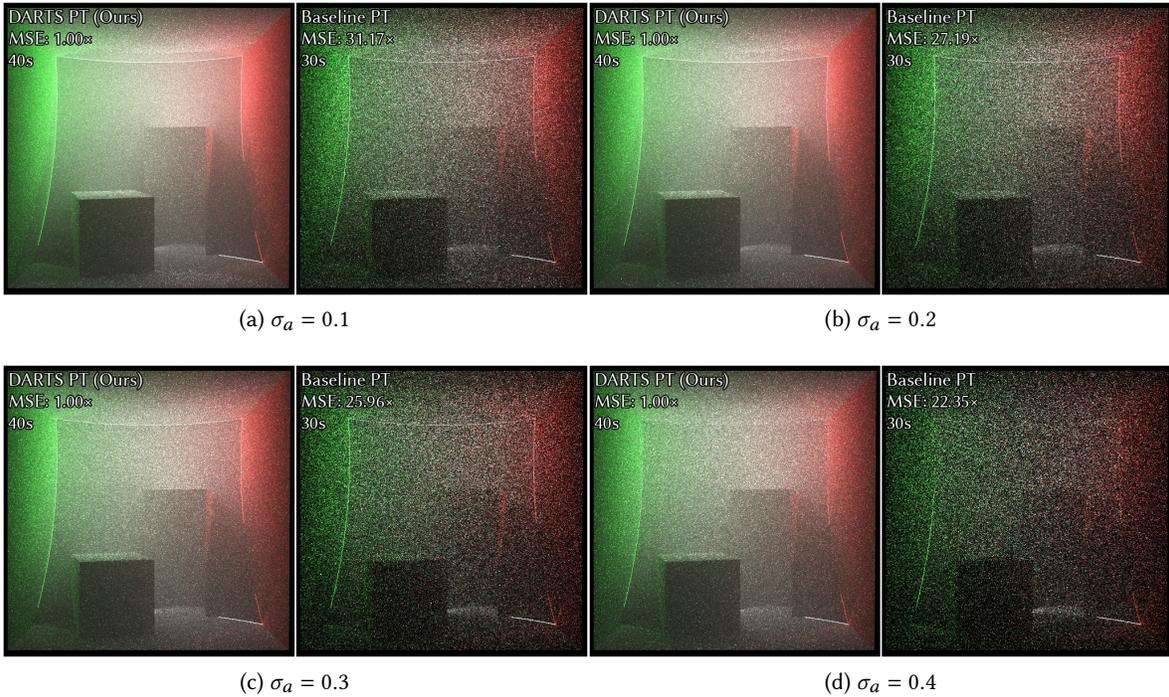


Fig. XI. Violation to the assumptions of diffusion approximation: absorption coefficients. In this experiments, four different σ_a s are presented and the amplitude of the parameters are not negligible compared to the σ_s (0.4).

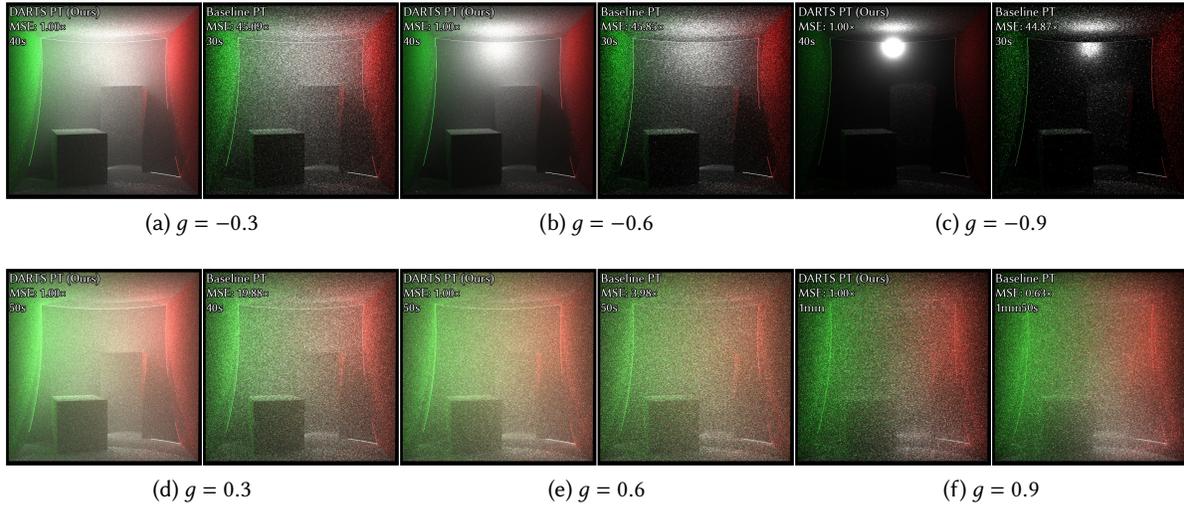


Fig. XII. Violation to the assumptions of diffusion approximation: scattering directionality (g in HG phase function). In these experiments, six different $\sigma_a s$ are presented. While for the backward scattering cases, our method remains robust, the performance in forward scattering cases degrades as g increases.

C.6 Peaky phase function cases

In this section, we briefly illustrate the sufficiency of phase function sampling for the MIS of EDA sampling. Since our elliptical sampling is a two-step approach and in the first step, we sample (or reuse) the direction of the elliptical connection, for media with strongly directed phase function, it seems theoretically inefficient for phase function sampling, due to the fact that on the ellipse, the throughput is determined by the product of two consecutive phase functions. Although the phase function requires being evaluated twice, these evaluations are not independent. The direction computed during the first evaluation determines the cosine value input for the phase function during the second evaluation. As a result, the shape of the product of these two phase functions is not significantly different from that of a single local phase function, even in cases where the phase function is peaky. As shown in Figure XIII, we illustrate the relative magnitudes of the function by tabulating directions originating from a specific focal point of the ellipse. We employ Monte Carlo methods to compute relative value plots of the function under different phase function parameters and strategies, with normalization such that the maximum value is 1. The first row corresponds to a single local phase function, while the second row corresponds to the product of two phase functions.

REFERENCES

- Benedikt Bitterli. 2016. Rendering resources. <https://benedikt-bitterli.me/resources/>.
- 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.
- PyTorch Contributors. 2023. PyTorch: An open source machine learning framework. <https://github.com/pytorch/pytorch/releases/tag/v2.0.0>. Version 2.0.0.
- 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).

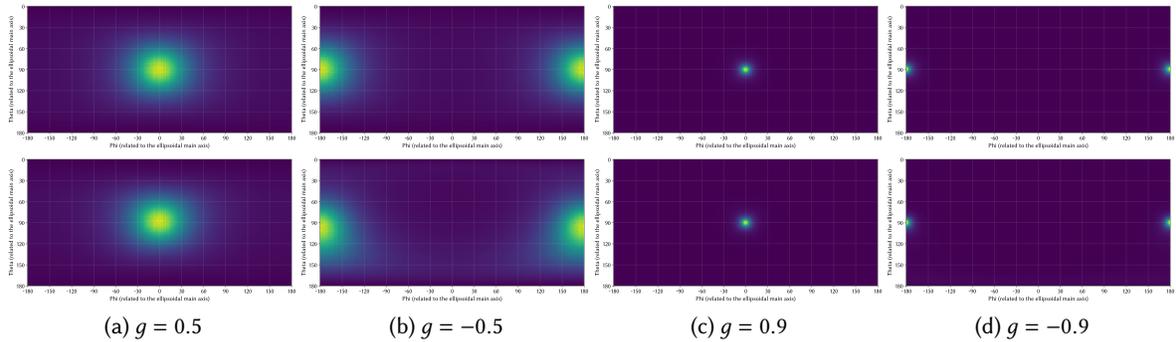


Fig. XIII. Sufficiency of phase function sampling for the MIS of EDA sampling. The first row depicts the relative value of single local phase function, while the second row depicts the relative value of the product of two phase functions. The columns varies in the directionality of phase function: $g \in [-1, 1]$ represents the parameter of Henyey-Greenstein phase function. Higher absolute value means the phase function is more directed (peaky).

- Adrian Jarabo, Julio Marco, Adolfo Munoz, Raul Buisan, Wojciech Jarosz, and Diego Gutierrez. 2014. A framework for transient rendering. *ACM Transactions on Graphics (ToG)* 33, 6 (2014), 1–10.
- 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.
- 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>
- Matt Pharr, Wenzel Jakob, and Greg Humphreys. 2023. *Physically based rendering: From theory to implementation*. MIT Press.