

Geometric Antithetic Sampling for Rendering Spatiotemporally Modulated Light

JUHYEON KIM, Dartmouth College, USA

JACK CUI, Dartmouth College, USA

WOJCIECH JAROSZ, Dartmouth College, USA

ADITHYA PEDIREDLA, Dartmouth College, USA




Fig. 1. We propose an efficient sampling technique for modulated light transport, including temporally modulated continuous-wave time-of-flight (CW-ToF) and spatially modulated structured light systems. The key idea is to form antithetic pairs by geometrically deforming the primal path such that the modulation terms have opposite signs. We demonstrate the usefulness of our method for reducing variance in MC rendering and in various downstream applications.

We propose a geometric antithetic sampling strategy for efficiently rendering spatiotemporally modulated light, encompassing continuous-wave time-of-flight (CW-ToF) and structured light systems. This is motivated by observing that modulated light transport can be decomposed into a slowly varying geometric term and a rapidly oscillating modulation term, where the latter dominates variance. Our key idea is to construct antithetic sample pairs by geometrically deforming a light path so that the modulation term changes sign while preserving similarity in the underlying transport. In both CW-ToF and structured light, antithetic samples lie on an antithetic manifold—the set of configurations with opposite modulation—and we search along it to find the path most similar to a given reference path, thereby maintaining strong negative correlation between paired samples. For CW-ToF, this corresponds to shifting paths in optical path-length space, while for structured light, it involves deforming emitter samples along the spatial modulation manifold in projector-coordinate space. By unifying both cases under this geometric deformation framework, our method cancels high-frequency modulation components through paired samples, leading to significant variance reduction and enabling more efficient simulation of coded illumination and sensing systems in computational imaging.

CCS Concepts: • **Computing methodologies** → **Ray tracing; Computational photography.**

Authors' Contact Information: Juhyeon Kim, Dartmouth College, Hanover, NH, USA, juhyeon.kim.gr@dartmouth.edu; Jack Cui, Dartmouth College, Hanover, NH, USA, jack.cui.gr@dartmouth.edu; Wojciech Jarosz, Dartmouth College, Hanover, NH, USA, wojciech.k.jarosz@dartmouth.edu; Adithya Pediredla, Dartmouth College, Hanover, NH, USA, adithya.k.pediredla@dartmouth.edu.



This work is licensed under a Creative Commons Attribution 4.0 International License. SA Conference Papers '26, Kuala Lumpur, Malaysia
© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2842-6/2026/12
<https://doi.org/10.1145/3829340.3842310>

Additional Key Words and Phrases: Structured Light, Continuous-wave Time-of-Flight Imaging, Modulated Light Rendering

ACM Reference Format:

Juhyeon Kim, Jack Cui, Wojciech Jarosz, and Adithya Pediredla. 2026. Geometric Antithetic Sampling for Rendering Spatiotemporally Modulated Light. In *SIGGRAPH Asia 2026 Conference Papers (SA Conference Papers '26)*, December 01–04, 2026, Kuala Lumpur, Malaysia. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3829340.3842310>

1 Introduction

Modulated light sources play a critical role in many computational imaging systems. By encoding information into the temporal or spatial variation of illumination, these systems enable the recovery of scene geometry and appearance beyond what is possible with conventional intensity measurements [Zanuttigh et al. 2016]. Two prominent examples are (1) continuous-wave time-of-flight (CW-ToF), which uses *temporally* modulated illumination to estimate time-of-flight encoded in the phase delay, and (2) structured light, which uses *spatially* modulated projection patterns to reconstruct scene geometry from their observed deformation (Fig. 1). In this work, we show that light transport in both CW-ToF and structured-light systems can be formulated within a unified framework in which the measured signal is expressed as a path integral augmented by a path-dependent modulation function. This path-space perspective reveals a shared structure between these otherwise distinct imaging modalities, suggesting a common strategy for their efficient simulation—the *antithetic* sampling approach we propose.

Accurately simulating light transport under modulated illumination is crucial for both system design and algorithm development. Many existing computational imaging methods rely on simplified

forward models that assume direct illumination. However, in real-world scenarios, global illumination effects such as interreflections, shadows, and indirect scattering significantly influence the measured signal. Ignoring these effects can lead to systematic errors in reconstruction algorithms and limit the performance of sensing systems. Therefore, a physically accurate rendering framework that accounts for full light transport under modulated illumination is highly desirable, as it enables a better understanding of measurement formation and facilitates the development of more robust and accurate reconstruction techniques.

A key challenge in simulating modulated light transport lies in the presence of a highly oscillatory modulation term. While the geometric component of light transport varies relatively smoothly, the modulation term often introduces rapid variations that dominate the variance of Monte Carlo (MC) estimators. Naive MC sampling fails to efficiently handle such oscillatory behavior, as independently sampled paths exhibit poor modulation correlation, leading to slow convergence. Addressing this challenge requires sampling strategies that explicitly exploit the structure of the modulation term.

In this work, we propose an efficient *antithetic sampling* strategy for modulated light transport. Our approach constructs correlated sample pairs by *geometrically* deforming an existing light path so that the associated modulation terms become antithetic—i.e., equal in magnitude but opposite in sign—while preserving similarity in the underlying geometric transport quantity. This effectively cancels the oscillatory component of the signal and significantly reduces variance. Our unified formulation enables efficient variance reduction across both temporal and spatial modulation settings, providing a practical framework for rendering various modulated light transport scenarios in computational imaging systems.

The code and data are available on the project page¹.

2 Related Works

2.1 Modulated Light in Computational Imaging

Modulated imaging systems can be broadly interpreted as employing either temporal or spatial modulation.

Continuous-wave time-of-flight (CW-ToF), also referred to as indirect Time-of-Flight (iToF), is a representative technique in computational imaging that uses temporally modulated illumination. It uses amplitude-modulated continuous-wave (AMCW) illumination and infers time-of-flight from the phase shift between the emitted and received signals, thereby providing scene depth [Hansard et al. 2012; Horaud et al. 2016]. Such systems have been widely adopted because of their high acquisition speed and compact hardware design, and have found broad applications in robotics [Alenyà et al. 2014; Foix et al. 2011], autonomous driving [Li and Ibanez-Guzman 2020], and consumer depth cameras [Izadi et al. 2011].

Structured light is a representative spatially modulated imaging technique in which coded illumination patterns are projected onto a scene and observed by a camera [Zanuttigh et al. 2016]. Scene geometry can then be reconstructed by analyzing the spatial deformation or correspondence of the projected patterns. Common structured light approaches include binary coding patterns such as Gray codes [Salvi et al. 2004] and sinusoidal fringe projection

profilometry (FPP) [Zuo et al. 2018]. Compared to passive stereo methods, structured light often provides higher reconstruction accuracy and robustness in textureless regions.

Both CW-ToF and structured light systems are typically based on a single-bounce imaging model and therefore suffer from artifacts caused by global illumination, commonly referred to as *multipath interference*. A common solution is to employ high-frequency modulation, where rapidly oscillating positive and negative modulation components tend to cancel out under global illumination [Nayar et al. 2006]. For example, some CW-ToF systems employ GHz-range modulation [Baek et al. 2023; Gupta et al. 2015], while some structured light methods use high-frequency spatial patterns to reduce indirect illumination effects [Gupta et al. 2011; Nayar et al. 2006]. Our antithetic sampling is particularly well suited for simulating such high-frequency modulated indirect transport, enabling accurate evaluation of subtle multipath effects as well as the analysis and design of new modulation codes.

2.2 Modulated Light Rendering

CW-ToF belongs to a family of ToF imaging systems, but many existing ToF rendering approaches focus on simulating either transient image sequences [Jarabo et al. 2014; Marco et al. 2017] or time-gated images within narrow temporal gates [Liu et al. 2022; Pediredla et al. 2019]. These formulations are not directly applicable to CW-ToF, which involves temporally modulated measurements. One of the most relevant works is that of Kim et al. [2023], who proposed efficient rendering methods for Doppler ToF cameras, a variant of CW-ToF cameras designed for velocity imaging [Heide et al. 2015]. They employ *temporal* antithetic sampling with path correlation to handle one-period modulation over the exposure time in dynamic scenes. In contrast, we assume static scenes and perform antithetic sampling by *geometrically* deforming paths to achieve opposite modulation. Another important inspiration comes from Kim et al. [2026], who introduce path-length-aware shift mapping to preserve path length during ReSTIR-style path reuse [Bitterli et al. 2020; Lin et al. 2022]. We also exploit this formulation for CW-ToF, but adapt the path length increment to satisfy the antithetic condition.

Structured light has been studied much less from a rendering perspective. Several works employed image-based rendering for live performance simulation [Jones et al. 2006] or physically based rendering to reproduce benchmark structured-light systems [Medeiros et al. 2014]. In contrast, our work focuses more on MC variance reduction by explicitly considering the underlying light transport.

2.3 Antithetic Sampling

Antithetic sampling is a variance reduction technique that exploits two negatively correlated samples [Hammersley and Morton 1956]. It has found various applications in rendering, including improved jittered sampling [Subr et al. 2014], differentiable rendering [Bangaru et al. 2020; Wu et al. 2021; Yu et al. 2022; Zhang et al. 2021], gradient-domain rendering [Kettunen et al. 2015; Manzi et al. 2016] and residual rendering [Xu et al. 2024]. To construct highly correlated sample pairs, these works employ specially designed path-mapping functions, often referred to as *shift mappings*, which can be implemented either in primary sample space [Hua et al. 2019;

¹https://juhyeonkim95.github.io/project-pages/antithetic_modulation

Manzi et al. 2016] or through vertex reconnection [Kettunen et al. 2015]. In our work, however, the modulation condition constrains the antithetic path sample to lie on a specific *antithetic manifold* (similar to the specular manifold [Jakob and Marschner 2012], but defined by the modulation value), preventing us from directly applying existing shift-mapping methods. Instead, we minimize the displacement toward the manifold to preserve path similarity, as in Kim et al. [2026], but our use of antithetic sampling enables a simplified MIS evaluation that was not possible in their setting.

3 Proposed Method

3.1 Modulated Light Transport

The standard path-integral formulation of light transport [Veach 1997] is

$$I = \int_{\mathcal{P}} f(\bar{x}) d\mu(\bar{x}), \quad (1)$$

where I denotes the measurement at a pixel, $\mathcal{P}$ is the space of all possible light paths $\bar{x} = (\mathbf{x}_0, \dots, \mathbf{x}_K)$, with $\mathbf{x}_0$ at the camera and $\mathbf{x}_K$ at the emitter, $f(\bar{x})$ is the path throughput containing the geometric, BSDF, emission, and visibility terms, and $d\mu(\bar{x})$ denotes the path-space measure.

Modulated light transport can similarly be expressed as

$$I = \int_{\mathcal{P}} f(\bar{x}) m(\bar{x}) d\mu(\bar{x}), \quad (2)$$

where $m(\bar{x})$ is a path-dependent *modulation* term. For CW-ToF, it depends on the path length, whereas for structured light, it depends on the projector coordinates.

Zero-mean Modulation. Any modulation m can be decomposed into a constant DC component m_{DC} and a zero-mean (over one period) oscillatory component $\tilde{m}$:

$$m(\bar{x}) = m_{\text{DC}} + \tilde{m}(\bar{x}). \quad (3)$$

Substituting this decomposition into Eq. (2) gives

$$I = m_{\text{DC}} \int_{\mathcal{P}} f(\bar{x}) d\mu(\bar{x}) + \int_{\mathcal{P}} f(\bar{x}) \tilde{m}(\bar{x}) d\mu(\bar{x}). \quad (4)$$

The first term is a scaled standard path integral and can therefore be estimated using conventional rendering methods. Our antithetic method targets the second term, particularly in high-frequency settings where positive and negative modulation values alternate rapidly over path space. The full physical measurement can be recovered by adding the DC term. Note that the unmodulated integral is independent of the modulation pattern, so it can be rendered once and shared across multiple patterns, as is common in modulated-light applications. For simplicity, we will use m to denote $\tilde{m}$ throughout the remainder of the paper.

3.2 Antithetic Sampling for Modulated Light Transport

Our key idea is to apply antithetic sampling, a technique widely used in the rendering literature (see Sec. 2.3), but to cancel opposing modulation contributions in the modulated path integral:

$$m(\bar{x}_p) = -m(\bar{x}_a), \quad f(\bar{x}_p) \approx f(\bar{x}_a) J_T, \quad (5)$$

where $\bar{x}_p$ is a primal path, $\bar{x}_a = T(\bar{x}_p)$ is an antithetic path via mapping $T : \mathcal{P} \rightarrow \mathcal{P}$ and $J_T = |d\bar{x}_a/d\bar{x}_p|$ is the Jacobian determinant

of the mapping. Because T operates directly in path space by geometrically deforming the primal path, we refer to our approach as *geometric antithetic sampling*. The opposite-modulation condition is enforced by construction, whereas the path throughput similarity is a variance-reduction objective rather than a requirement for unbiasedness. When the approximation is accurate, the paired contributions largely cancel, reducing variance. Unbiasedness is instead ensured by the MIS estimator introduced in the following section, even when path similarity is poor or T does not exist.

Our overall strategy is to find a path with exactly opposite modulation while preserving the path throughput, as described in Eq. (5). In fact, enforcing a particular modulation value defines a thin manifold in the path space, $\mathcal{M}(l) = \{\bar{x} \mid m(\bar{x}) = l\}$, which we refer to as the *modulation manifold* (red lines in Fig. 2). Given a primal sample $\bar{x}_p$, the condition $m(\bar{x}_a) = -m(\bar{x}_p)$ similarly defines a manifold of possible antithetic paths $\bar{x}_a$, which we call the *antithetic manifold* (dotted red line crossing $\mathbf{p}'$ in Fig. 2). In general, finding a point on the antithetic manifold that also preserves path throughput is highly challenging. Instead, we seek a mapping T that reaches the antithetic manifold while minimizing geometric displacement, based on the intuition that smaller path deformation leads to more similar path throughput.

Antithetic MIS Estimator. Exploiting multiple importance sampling (MIS) between primal and antithetic samples yields

$$\hat{I} = w_p^{\text{MIS}}(\bar{x}_p) \frac{f(\bar{x}_p) m(\bar{x}_p)}{p_p(\bar{x}_p)} + w_a^{\text{MIS}}(\bar{x}_a) \frac{f(\bar{x}_a) m(\bar{x}_a)}{p_a(\bar{x}_a)}. \quad (6)$$

Each MIS weight w^{MIS} from the balance heuristic is given by

$$w_p^{\text{MIS}}(\bar{x}_p) = \frac{p_p(\bar{x}_p)}{p_p(\bar{x}_p) + p_a(\bar{x}_p)}, \quad w_a^{\text{MIS}}(\bar{x}_a) = \frac{p_a(\bar{x}_a)}{p_a(\bar{x}_a) + p_p(\bar{x}_a)}, \quad (7)$$

where p_p is the PDF of the primal sample and p_a is the PDF of the antithetic sample, given by $p_a(\bar{x}) = p_p(T^{-1}(\bar{x}))/J_T$. In fact, MIS between mapped paths is important not only for reducing variance but also for guaranteeing unbiasedness when the mapping fails (e.g., visibility change or discontinuity), which is already known in the shift-mapping literature (see Lin et al. [2022] for details). In this case, the antithetic proposal has zero density, yielding $w_p^{\text{MIS}} = 1$ and $w_a^{\text{MIS}} = 0$; the estimator therefore falls back to the primal estimator without introducing bias.

Simplified MIS Weight Evaluation. One caveat in evaluating MIS weights is that to obtain $p_a(\bar{x}_p)$, we need to evaluate $T^{-1}(\bar{x}_p)$, which corresponds to another sample distinct from $\bar{x}_p$ and $\bar{x}_a$. However, if the mapping is *involutory*, i.e., $T(T(\bar{x}_p)) = \bar{x}_p$, then $T^{-1}(\bar{x}_p) = \bar{x}_a$, so we do not need to evaluate an additional mapping. Therefore, it is desirable for the mapping function to be involutory, as this effectively halves the computational cost of performing the antithetic mapping. More generally, if T is stochastically chosen from two transforms $\{T_1, T_2\}$ satisfying the inverse relation $T_1^{-1} = T_2$ (e.g., path-length increase and decrease) with equal probability, MIS weights can be evaluated without computing an additional mapping. The derivation is provided in the supplementary material. It is worth noting that this MIS simplification generally does not apply to standard shift-mapping methods including Kim et al. [2026], which operate between distinct base and offset path domains, $T : \mathcal{P} \rightarrow \mathcal{P}'$.




Fig. 2. Geometric construction of antithetic sample pairs for two imaging modalities. (A) For CW-ToF, the antithetic manifold consists of paths whose lengths differ by half a wavelength ($\delta\ell = \pm\lambda/2$). An average-gradient constraint selects a geometrically similar antithetic path while maintaining bijectivity. (B) For structured light, the antithetic manifold is determined directly by projector coordinates with opposite pattern values.

Even if T and T^{-1} are bijective, the distinction between the domains must be retained when evaluating the corresponding densities. On the other hand, in our antithetic setting, both samples belong to the same path domain, $T : \mathcal{P} \rightarrow \mathcal{P}$, enabling the simplified MIS evaluation.

The actual antithetic strategy depends on the imaging modality. In the following sections, we discuss antithetic sampling for CW-ToF (temporally modulated) and structured light (spatially modulated).

3.3 Antithetic Sampling for CW-ToF

CW-ToF is a time-of-flight imaging system in which the modulation term depends on a path-length importance function W , defined by the correlation of two temporal modulation signals:

$$m(\bar{\mathbf{x}}) = W(\ell(\bar{\mathbf{x}})) := \int_0^{T_{\text{exp}}} s(t)g(t - \ell(\bar{\mathbf{x}})/c) dt, \quad (8)$$

where $s(t)$ and $g(t)$ denote the modulation functions at the camera and emitter sides, respectively, and T_{exp} is the exposure time. The optical path length $\ell(\bar{\mathbf{x}})$ is given by $\sum_{k=0}^{K-1} \eta_k \|\mathbf{x}_{k+1} - \mathbf{x}_k\|$ where η_k is the refractive index of the path segment $\mathbf{x}_k \mathbf{x}_{k+1}$. Throughout, we use units in which the speed of light is $c = 1$, allowing optical path length and travel time to be used interchangeably, as is common in previous works [Kim et al. 2025, 2023; Pediredla et al. 2019].

For conventional CW-ToF systems, the modulations are sinusoidal, $s(t) = \cos(\omega t + \psi)$, $g(t) = \cos(\omega t)$, where ω is the modulation frequency and ψ is a phase offset. Substituting these modulation functions yields

$$s(t)g(t - \ell(\bar{\mathbf{x}})) = \frac{1}{2} [\cos(\omega\ell(\bar{\mathbf{x}}) + \psi) + \cos(2\omega t - \omega\ell(\bar{\mathbf{x}}) + \psi)]. \quad (9)$$

The latter high-frequency term vanishes after integration over T_{exp} since $T_{\text{exp}} \gg 1/\omega$. Therefore, we typically consider only the first term, resulting in

$$m(\bar{\mathbf{x}}) = \cos(\omega\ell(\bar{\mathbf{x}}) + \psi), \quad (10)$$

where we omit the scale factor of $T_{\text{exp}}/2$, which is a constant term.

Antithetic Sampling Strategy. Since the modulation term is given by a sinusoidal function of the path length, our goal is to find an antithetic path whose path length differs by half a wavelength, $\lambda/2$,

where $\lambda = 2\pi/\omega$, such that the modulation terms have opposite signs². To achieve this, we exploit the *path-length-aware shift mapping* proposed by Kim et al. [2026], which was originally introduced to preserve path length after ReSTIR-style path reuse in time-gated rendering. To elaborate, we perturb a single path vertex $\mathbf{p}$ in the primal path and find an antithetic sample $\mathbf{p}' = T(\mathbf{p})$ satisfying

$$\ell(\mathbf{p}') = \ell(\mathbf{p}) + \delta\ell, \quad (11)$$

where $\ell(\mathbf{p}) = |\mathbf{p}_1 - \mathbf{p}| + |\mathbf{p}_2 - \mathbf{p}|$ is the sum of distances to the neighboring path vertices $\mathbf{p}_1$ and $\mathbf{p}_2$, and $\delta\ell = \pm\lambda/2$ is stochastically selected with equal probability. $\mathbf{p}$ can be any path vertex, but for simplicity, we limit $\mathbf{p}$ to the NEE vertex (i.e., the vertex immediately preceding the emitter vertex), where $\mathbf{p}_2 = \mathbf{x}_K$ and $\mathbf{p} = \mathbf{x}_{K-1}$. Using a 2D coordinate parameterization of $\mathbf{p}$ with $\xi \in \mathbb{R}^2$ (e.g., barycentric coordinates or hemisphere coordinates from a neighboring path vertex), this condition can be simplified to

$$\ell(\xi') = \ell(\xi) + \delta\ell. \quad (12)$$

The unknown variable ξ' is two-dimensional, whereas the path-length condition imposes only a one-dimensional constraint. As a result, the condition defines only an antithetic manifold, as illustrated in Fig. 2-(A). To resolve this ambiguity, we search along the average-gradient direction $\nabla\ell_{\text{avg}} = (\nabla\ell(\xi) + \nabla\ell(\xi'))/2$; this minimizes the displacement while guaranteeing bijectivity. This gives a 2D constraint F :

$$F(\xi, \xi') = \begin{bmatrix} \ell(\xi') - \ell(\xi) - \delta\ell \\ (R\nabla\ell_{\text{avg}}) \cdot (\xi' - \xi) \end{bmatrix} = \begin{bmatrix} 0 \\ 0 \end{bmatrix}, \quad (13)$$

where R is a 90-degree rotation matrix.

We solve for ξ' using Newton's method, initializing it at ξ and iteratively updating it as follows:

$$\xi'_0 = \xi, \quad \xi'_{k+1} \leftarrow \xi'_k - \left(\frac{\partial F}{\partial \xi'_k} \right)^{-1} F(\xi, \xi'_k). \quad (14)$$

We omit implementation details, such as the evaluation of the derivative of F and the mapping Jacobian, and refer readers to the supplementary material and Kim et al. [2026] for further details.

3.4 Antithetic Sampling for Structured Light

For structured light, the modulation term depends not on the path length but on the projector coordinates:

$$m(\bar{\mathbf{x}}) = P(\xi(\bar{\mathbf{x}})), \quad (15)$$

where $P : [0, 1]^2 \rightarrow \mathbb{R}$ is the projector pattern and $\xi = (\xi_1, \xi_2) \in [0, 1]^2$ represents the projector coordinates determined by the emitter-connecting path segment $\mathbf{x}_{K-1} \mathbf{x}_K$. We keep the notation ξ consistent with the CW-ToF formulation to show the similarity between the two objectives. P can be an arbitrary function, but in this work we focus on two commonly used cases: 1D sinusoidal and binary patterns (Fig. 2-(B)), which are widely used in FPP and coded illumination, respectively. Specifically, they can be written as

$$P(\xi) = \cos(v), \quad P(\xi) = 2B(\lfloor v \rfloor) - 1, \quad (16)$$

where $v = \xi_1/\lambda$ (or ξ_2/λ), and B denotes a binary code (0s and 1s) of length $N = 1/\lambda$, which we assume to be periodic (e.g., a Gray code).

²One can also use other periodic modulations (e.g., triangular, square, sawtooth), but these may require distinct mappings to produce opposite modulation values (see Fig. 6).




Fig. 3. Antithetic mappings for binary structured-light codes. (A) Periodic codes and (B) aperiodic codes with antithetic blocks admit direct pairing. For arbitrary codes without such blocks, we use (C) optimal transport, with indices of 0s and 1s indicated, or (D) scale-based mapping between contiguous blocks, with the corresponding nonuniform Jacobians shown.

Antithetic Sampling Strategy. Our goal is to find an antithetic projector coordinate $\xi' = T(\xi)$ that satisfies:

$$P(\xi') = -P(\xi). \quad (17)$$

This is straightforward for structured light since the antithetic manifold is directly defined by the projector pattern³. As in CW-ToF, there is still an additional degree of freedom, but selecting a point on this manifold is simple: we fix one axis and shift the other axis by half a wavelength, i.e., $(\xi_1, \xi_2) \rightarrow (\xi_1 \pm \lambda/2, \xi_2)$. This construction minimizes displacement in ξ -space while preserving bijectivity. We omit the details for 2D patterns (e.g., 2D sinusoidal patterns or checkerboards), but the same antithetic strategy can be applied.

Above, we assumed periodic modulation, but in practice, modulation need not be periodic, particularly for binary codes. Many systems employ aperiodic codes composed of irregular arrangements of 0s and 1s. Interestingly, antithetic sampling can still be effectively applied even in the absence of periodicity.

3.4.1 Aperiodic Codes with Antithetic Blocks. We first consider a simpler case in which the code can be partitioned into multiple *antithetic blocks*, each containing equal numbers of 0s and 1s, as illustrated in Fig. 3-(B). In this setting, the antithetic mapping can be defined locally within each block rather than through a globally uniform rule.

One example is the XOR code proposed by Gupta et al. [2011]. For the XOR-02 code (Fig. 8), constructed by taking the bitwise XOR of a Gray-code pattern with the highest-frequency pattern, antithetic blocks of size 2 always exist. Similarly, XOR-04 always admits antithetic blocks of size 4. More generally, whenever a code is composed of stacked antithetic blocks, we can exploit this local structure to perform antithetic sampling.

3.4.2 Arbitrary Aperiodic Codes. A more challenging case arises when such antithetic blocks cannot be identified, as shown in Fig. 3-(C, D). Even in this setting, it is still possible to construct antithetic mappings between code elements.

Optimal Transport-based Mapping. The first strategy exploits optimal transport (OT)-based mapping. OT-based mapping minimizes the total displacement of the mapping, thereby preserving path similarity as much as possible on average. More specifically, optimal transport defines a mapping that minimizes the transport distance between the code and its antithetic counterpart:

$$T_{OT} = \arg \min_T \sum_{i=0}^{N-1} |i - T(i)|, \quad (18)$$

where $i = \lfloor v \rfloor$ is the code index and $T(i)$ is the antithetic mapping satisfying $B(i) \neq B(T(i))$. This can also be interpreted as a transport problem between the two discrete distributions $\mathbb{1}_{B(i)=0}$ and $\mathbb{1}_{B(i)=1}$. The minimum transport cost is known as the *Earth Mover's Distance* (EMD), and the optimal transport solution directly yields this value. The effectiveness of the OT-based method depends on the EMD of the code. When the EMD is sufficiently small, path similarity is well preserved. In contrast, large EMD values lead to path throughput mismatch, analogous to the degradation observed when using large wavelengths in the periodic case.

The implementation of OT-based mapping is straightforward. Scanning from left to right, we index the 0s and 1s independently and match the k -th 0 to the k -th 1, as illustrated in Fig. 3-(C).

Scale-based Mapping. Another possible strategy is scale-based mapping, which maps contiguous blocks of 0s and 1s, as shown in Fig. 3-(D). This method better preserves locality, but because the interval sizes are generally unequal, the Jacobian determinant $J_{T,\xi} = |\partial \xi' / \partial \xi|$ is no longer equal to 1, and therefore the cancellation is not perfect. A detailed comparison of the mapping strategies is provided in the results section (Sec. 4.4).

4 Results

We implement the proposed algorithm on top of the FalcOR framework [Kallweit et al. 2022]. We assume a point light source for CW-ToF and a pinhole projector for structured light. For most figures, we visualize only indirect illumination, since direct illumination is trivial under our setting. Additionally, because the modulation is zero-mean, we use a diverging luminance colormap to better highlight the oscillatory global-illumination components. All experiments are compared using equal time budgets.

4.1 Rendering Comparison

Wavelength Comparison. We first compare the rendering performance of the proposed antithetic sampling method against naive MC sampling. Specifically, we evaluate the mean squared error (MSE) of both CW-ToF and structured light rendering under sinusoidal modulation with different wavelengths. The results are shown in Fig. 4, with corresponding rendered images provided in Fig. 10. Compared to naive sampling, antithetic sampling is significantly more effective at canceling high-frequency global modulation. However, its effectiveness decreases as the modulation wavelength increases. This occurs because longer spatial wavelengths require larger displacements to construct antithetic pairs, reducing path similarity. For CW-ToF, larger displacements also increase the failure rate of Newton's method, as explained further in the following paragraph. We empirically find that antithetic sampling is more effective when

³For binary codes, we additionally consider the fractional component to maintain a continuous manifold.




Fig. 4. Equal-time MSE comparison of naive and antithetic sampling across three scenes (0.3, 1.2, and 1.2 s, respectively) and modulation wavelengths. Antithetic sampling performs best at shorter wavelengths, where paired paths remain geometrically similar, but its advantage decreases as wavelength increases. Corresponding images are shown in Fig. 10.

the modulation wavelength is below a threshold of approximately 1% – 10% of the scene scale.

Newton Iteration Statistics for CW-ToF. We report statistics of the Newton solver for CW-ToF across different wavelengths in Fig. 5, using the experiments shown in Fig. 4. As the wavelength increases, the required path deformation becomes larger, reducing the success rate and increasing the average number of iterations. Nevertheless, successful mappings typically converge within two to three iterations, using a convergence tolerance of 0.001 λ .

For the 2D parameterization, we use local barycentric coordinates, which work well in most scenes. However, ray-tracing-based global coordinates may be preferable for scenes composed only of tiny primitives, where local coordinates frequently cross boundaries and cause mapping failures. Further details are provided in the supplementary material.




Fig. 5. Newton solver success rate and average iteration count for CW-ToF in Fig. 4. The black line shows the success rate, while the red line shows the average number of iterations. Longer wavelengths require larger path deformations, resulting in lower success rates and more iterations. Successful mappings generally converge within two to three iterations.

Computational Overhead. We report the computational overhead of antithetic sampling using the proposed mapping in Tab. 1. We measure the average SPP achieved under the same time budget in Fig. 4. The overhead varies with scene complexity, ranging from 65.4% to 150.7% for CW-ToF and from 32.2% to 91.1% for structured

Table 1. Average SPP and overhead across the experiments in Fig. 4.

Scene (SPP)	CW-ToF			Structured Light		
	Naive	Anti.	Overhead	Naive	Anti.	Overhead
CORNELL BOX	302.9	183.3	65.4%	292.4	225.0	32.2%
STAIRCASE	539.0	268.4	101.1%	575.1	377.0	57.6%
KITCHEN	515.5	206.8	150.7%	514.9	284.0	91.1%

light. CW-ToF has higher overhead because constructing an antithetic path requires an iterative process, whereas structured-light mappings are evaluated directly in projector coordinates.

Waveform Comparison. We compare different modulation waveforms, including triangular, square, and sawtooth waves, for structured light with $\lambda = 0.01$ in Fig. 6. For triangular and square waves, which we define to be symmetric about the center of each period, translating by $\lambda/2$ produces an antithetic pair $(\xi, \xi + 0.5\lambda)$, similar to sinusoidal modulation. In contrast, such a translation does not produce perfectly matching antithetic pairs for sawtooth waves; instead, a mirror mapping $(\xi, \lambda - \xi)$ is required to obtain opposite modulation (we omit the modulo- λ for simplicity). This mirror mapping does not work for the triangular or square waves defined above. However, if the waveform is reparameterized with a $\lambda/4$ phase offset, the mirror mapping also becomes valid, while translation remains valid. Overall, the construction of valid antithetic pairs depends on both the waveform and the definition of the period, while translating and mirroring serve as the two fundamental construction strategies.




Fig. 6. Comparison of antithetic mappings for triangular, square, and sawtooth patterns. Translation pairs $(\xi, \xi + \lambda/2)$, while mirroring pairs $(\xi, \lambda - \xi)$ modulo λ . Translation is valid for triangular and square waves, whereas mirroring is required for the sawtooth and also becomes valid for the others after a $\lambda/4$ phase shift.

MIS Comparison. We also compare different MIS strategies in Tab. 2. Without MIS, the rendered results exhibit higher noise and often suffer from fireflies. In contrast, our proposed MIS formulation achieves better overall performance. We additionally implement a non-bijective MIS variant that explicitly recomputes the inverse mapping instead of leveraging bijectivity. This additional computation reduces the achievable SPP by approximately 10 – 20%, resulting in higher MSE and demonstrating the effectiveness of our bijective MIS formulation as discussed in Sec. 3.2.

Table 2. Equal-time comparison of MIS strategies for CW-ToF with $\lambda = 0.02$. Each entry reports rendering MSE followed by the SPP in parentheses.

Scene / MSE (SPP)	w/o MIS	w/o bijectivity	full MIS
CORNELL-BOX	3.33 (169)	3.50 (144)	3.08 (165)
STAIRCASE	3.37 (253)	3.64 (208)	3.26 (238)
KITCHEN	7.69 (189)	5.98 (170)	5.31 (185)

Comparison with Ellipsoidal Connection. We also compare our method against ellipsoidal connections [Pediredla et al. 2019], which can similarly generate opposite-modulation antithetic samples by setting the focal distance sum to $\ell \pm \lambda/2$ (Fig. 7-(A)). However, this approach searches over all scene triangles and is agnostic to path correlation, making it ineffective for modulation cancellation. This issue can be partially alleviated by restricting the search to the same triangle (Fig. 7-(B)), but the results remain noisy, whereas our local search method achieves the best performance.




Fig. 7. Comparison of antithetic path construction strategies on STAIRCASE. (A) Ellipsoidal connection satisfies the path-length constraint but lacks path similarity. (B) Restricting it to the same triangle improves similarity. (C) Our local search best preserves path similarity and achieves the lowest MSE.

4.2 Depth Reconstruction with Structured Light

We demonstrate structured-light geometry reconstruction, for which rendering quality directly affects the accuracy of reconstruction from synthetic measurements.

Coded Illumination. We first perform depth reconstruction using coded illumination. The overall scene configuration is shown in Fig. 8-(A). A standard approach employs Gray code patterns to uniquely encode projector coordinates. However, in the presence of indirect illumination, the captured intensities no longer correspond to purely direct illumination, leading to decoding errors, as highlighted by the red arrows in the reconstruction results. To mitigate this issue, high-frequency coding strategies are used to suppress low-frequency global illumination. In particular, here we simulate the XOR-based codes of Gupta et al. [2011], examples of which are shown in Fig. 8-(A). We render these XOR patterns using both naive and antithetic sampling (following Sec. 3.4.1). We use 10-bit codes ($N = 1024$), together with an all-white reference image for normalization.

The first two rows of Fig. 8-(B) show rendering and reconstruction results using XOR-02. Naive sampling fails to produce clean reconstructions due to high variance in the estimated indirect illumination, which prevents reliable cancellation across XOR patterns. In contrast, our method significantly improves reconstruction quality by reducing variance in the indirect component and enabling stable modulation cancellation.




Fig. 8. (A) Depth reconstruction using structured light. Conventional Gray-code decoding often fails in the presence of indirect illumination. The high-frequency-only XOR patterns proposed by Gupta et al. [2011] can be used instead. (B) We render the XOR-02 measurements using naive and antithetic sampling and then perform depth reconstruction. Following Gupta et al. [2011], we also perform ensemble averaging over two XOR patterns to improve the reconstruction.

However, XOR-02 alone is insufficient to fully resolve the ambiguities. Even with GT rendering, artifacts remain due to code ambiguity within a single pixel and imperfect cancellation of indirect components. To address this issue, we additionally perform ensemble averaging with XOR-04 following Gupta et al. [2011]. The combination weights are determined from decoding confidence, measured using bit ambiguity and projection validity. This combined strategy further improves reconstruction quality. Nevertheless, noticeable noise artifacts remain with naive sampling, whereas our method produces clean and stable depth reconstructions.

Sinusoidal Illumination. We also demonstrate depth reconstruction using fringe projection profilometry (FPP) with sinusoidal projection patterns, which are widely used in applications such as automatic optical inspection (AOI) [Abd Al Rahman and Mousavi 2020]. We employ phase-shifting profilometry (PSP) to estimate the phase deformation induced by the object relative to a reference plane, followed by spatial phase unwrapping to resolve phase ambiguity. The reconstruction result for a circuit board mimicking an AOI scenario is shown in Fig. 11, using a sinusoidal pattern with $\lambda = 0.01$. Similar to coded illumination, we find that phase estimation is highly sensitive to rendering noise, and the high level of MC noise produced by naive sampling greatly degrades the reconstruction quality. In contrast, our antithetic sampling greatly improves phase stability and produces cleaner and more accurate object reconstructions.

We can also perform depth reconstruction for CW-ToF using sinusoidal modulation; we defer these results to the supplementary material.

4.3 Illumination Decomposition with Structured Light

We perform direct-global illumination decomposition using high-frequency checkerboard illumination following Nayar et al. [2006], as shown in Fig. 12-(A). The scene contains three objects with different material properties: diffuse, translucent, and specular. We use 25 shifted checkerboard patterns and perform decomposition by analyzing the minimum and maximum responses, following the original method. We visualize the error maps for each sampling strategy in Fig. 12-(B) (note that the error maps are identical for the direct and indirect components since they are complementary). Again, antithetic sampling achieves lower error than naive sampling, resulting in more accurate decomposition of direct and global illumination components.

4.4 Demonstration with Arbitrary Aperiodic Codes

We perform antithetic sampling using arbitrary binary codes of length 1024, generated by randomly shuffling equal numbers of 0s and 1s. We generate 50 different patterns and render each using naive sampling and antithetic sampling with both OT-based and scale-based mappings, as described in Sec. 3.4.2. We also compute the average EMD of each pattern and plot the MSE as a function of EMD across different patterns in Fig. 13-(A).

We observe that both antithetic sampling methods outperform naive sampling, which gives constant high MSE regardless of code. For the OT-based method, the rendering quality is strongly correlated with the EMD of the pattern, and we observe an approximately proportional relationship between EMD and MSE. In contrast, the performance of the scale-based method does not exhibit a clear dependence on EMD. Empirically, when the EMD is below 20 (2% of the total code length), the OT-based mapping performs better, while for larger EMD values, the scale-based method becomes more effective. The performance of the scale-based method is governed more by the mapping Jacobian than by the EMD. We plot the average Jacobian magnitude (using its reciprocal when < 1) against MSE for the scale-based method in the middle of Fig. 13-(A), and observe that more uniform Jacobians lead to better performance.

We visualize the rendered results for two representative patterns with different EMD values in Fig. 13-(B). For a small EMD, the OT-based method yields better results, whereas for a large EMD, the scale-based method performs better. We additionally show portions of the patterns and their corresponding OT mappings. When the EMD is small, the mapping remains close to a diagonal trend, indicating minimal displacement. In contrast, for large EMD, the mapping deviates a lot from the diagonal, indicating large transport distances. Overall, the OT- and scale-based strategies perform differently depending on the code EMD and average Jacobian, but both still offer an advantage over naive sampling.

Comparison with Stratified Sampling. Instead of using antithetic pairs, stratified sampling can be considered as an alternative approach for canceling modulation for arbitrary codes (i.e., assign one sample to each code element). Without path correlation, stratified sampling fails to achieve effective cancellation due to path mismatch. This can be improved by enforcing path matching (e.g., using the same ξ_2 coordinate), but doing so drastically reduces path diversity. We compare these methods in Fig. 9. Overall, our antithetic sampling




Fig. 9. Comparison with stratified sampling for arbitrary structured-light codes. Uncorrelated samples lack geometric similarity, while sharing ξ_2 reduces path diversity. Our antithetic mapping enables effective cancellation while preserving path diversity.

approach achieves the best performance by enabling modulation cancellation while preserving path diversity.

5 Conclusion

In this paper, we propose an efficient antithetic sampling strategy that geometrically deforms the original path to construct an antithetic path with opposite modulation while preserving path similarity. We demonstrate the effectiveness of the proposed method for both CW-ToF and structured light systems, in which antithetic sampling efficiently cancels highly oscillatory modulation components in global illumination. Our renderer enables efficient simulation of global illumination under complex high-frequency modulation, which is critical for accurate depth and geometry reconstruction in many computational imaging systems.

Future Work. First, we determine path similarity based on the parameter-space distance ($|\xi' - \xi|$) and optimize this by searching along the gradient direction in ξ -space. However, this does not always guarantee true path similarity. A more robust approach could incorporate auxiliary buffers, such as depth, normal, and albedo, rendered from the point-light or projector viewpoint, to better identify similar paths on the antithetic manifold.

We considered coded illumination only for structured light, but coding can also be applied to CW-ToF [Gupta et al. 2018; Kadambi et al. 2013], providing another direction for future work.

For arbitrary aperiodic codes, we explored both OT-based and scale-based mappings, each of which has its own advantages and limitations. It would be interesting to develop a unified strategy that performs well for both small- and large-EMD patterns. Furthermore, although we focused only on binary codes, similar ideas could potentially be extended to general continuous signals, for example, through optimal transport between their positive and negative components.

Finally, extending our framework to dynamic scenes would enable efficient simulation of motion-dependent effects, including motion blur in CW-ToF measurements.

Acknowledgments

We thank colleagues in the RISC and VCL labs for their helpful feedback. This work was supported in part by the National Science Foundation under Awards 2403122 and 2326904.

References

- M Abd Al Rahman and Alireza Mousavi. 2020. A review and analysis of automatic optical inspection and quality monitoring methods in electronics industry. *Ieee Access* 8 (2020), 183192–183271.
- Guillem Alenyà, Sergi Foix, and Carme Torras. 2014. ToF cameras for active vision in robotics. *Sensors and Actuators A: Physical* 218 (2014), 10–22.
- Seung-Hwan Baek, Noah Walsh, Ilya Chugunov, Zheng Shi, and Felix Heide. 2023. Centimeter-Wave Free-Space Neural Time-of-Flight Imaging. *ACM Transactions on Graphics* 42, 1 (March 2023), 3:1–3:18. doi:10/gspq49
- Sai Praveen Bangaru, Tzu-Mao Li, and Frédo Durand. 2020. Unbiased Warped-Area Sampling for Differentiable Rendering. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 39, 6 (Nov. 2020), 245:1–245:18. doi:10/gj32sk
- 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 (Proceedings of SIGGRAPH)* 39, 4 (July 2020). doi:10/gg8xc7
- Ji Deng, Jian Li, Hao Feng, Shumeng Ding, Yu Xiao, Wenzhong Han, and Zhoumo Zeng. 2020. Efficient intensity-based fringe projection profilometry method resistant to global illumination. *Optics Express* 28, 24 (2020), 36346–36360.
- Sergi Foix, Guillem Alenyà, and Carme Torras. 2011. Lock-in time-of-flight (ToF) cameras: A survey. *IEEE Sensors Journal* 11, 9 (2011), 1917–1926.
- Mohit Gupta, Amit Agrawal, Ashok Veeraraghavan, and Srinivasa G Narasimhan. 2011. Structured light 3D scanning in the presence of global illumination. In *CVPR 2011*. IEEE, 713–720.
- Mohit Gupta, Shree K Nayar, Matthias B Hullin, and Jaime Martin. 2015. Phasor Imaging: A Generalization of Correlation-Based Time-of-Flight Imaging. *ACM Transactions on Graphics* 34, 5 (2015). doi:10/gfz5k7
- Mohit Gupta, Andreas Velten, Shree K. Nayar, and Eric Breibach. 2018. What Are Optimal Coding Functions for Time-of-Flight Imaging? *ACM Transactions on Graphics* 37, 2 (April 2018), 1–18. doi:10/gd52p8
- John Michael Hammersley and Keith William Morton. 1956. A New Monte Carlo Technique: Antithetic Variates. *Mathematical Proceedings of the Cambridge Philosophical Society* 52, 03 (July 1956), 449–475. doi:10/dshxdn
- Miles Hansard, Seungkyu Lee, Ouk Choi, and Radu Patrice Horaud. 2012. *Time-of-Flight Cameras: Principles, Methods and Applications*. Springer-Verlag.
- Felix Heide, Wolfgang Heidrich, Matthias Hullin, and Gordon Wetzstein. 2015. Doppler Time-of-Flight Imaging. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 34, 4 (2015). doi:10/gfz5k9
- Radu Horaud, Miles Hansard, Georgios Evangelidis, and Clément Ménéier. 2016. An overview of depth cameras and range scanners based on time-of-flight technologies. *Machine vision and applications* 27, 7 (2016), 1005–1020.
- Changping Hu, Shuntao Liu, Dan Wu, and Jing Xu. 2023. Phase error model and compensation method for reflectivity and distance discontinuities in fringe projection profilometry. *Optics Express* 31, 3 (2023), 4405–4422.
- Binh-Son Hua, Adrien Gruson, Victor Petitjean, Matthias Zwicker, Derek Nowrouzezahrai, Elmar Eisemann, and Toshiya Hachisuka. 2019. A Survey on Gradient-Domain Rendering. *Computer Graphics Forum (Proceedings of Eurographics State of the Art Reports)* 38, 2 (2019), 455–472. doi:10/ggd8m5
- Shahram Izadi, David Kim, Otmár Hilliges, David Molyneaux, Richard Newcombe, Pushmeet Kohli, Jamie Shotton, Steve Hodges, Dustin Freeman, Andrew Davison, et al. 2011. Kinectfusion: real-time 3d reconstruction and interaction using a moving depth camera. In *Proceedings of the 24th annual ACM symposium on User interface software and technology*. 559–568.
- Wenzel Jakob and Steve Marschner. 2012. Manifold Exploration: A Markov Chain Monte Carlo Technique for Rendering Scenes with Difficult Specular Transport. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 31, 4 (July 2012), 58:1–58:13. doi:10/gfzq4p
- Adrian Jarabo, Julio Marco, Adolfo Munoz, Raul Buisan, Wojciech Jarosz, and Diego Gutierrez. 2014. A Framework for Transient Rendering. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 33, 6 (Nov. 2014), 177:1–177:10. doi:10/gfznb8
- Andrew Jones, Andrew Gardner, Mark Bolas, Ian McDowall, and Paul Debevec. 2006. Simulating spatially varying lighting on a live performance. In *The 3rd European Conference on Visual Media Production (CVMP 2006)-Part of the 2nd Multimedia Conference 2006*. IET, 127–133.
- A. Kadambi, R. Whyte, A. Bhandari, L. Streeter, C. Barsi, A. Dorrington, and R. Raskar. 2013. Coded Time of Flight Cameras: Sparse Deconvolution to Address Multipath Interference and Recover Time Profiles. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 32, 6 (2013). doi:10/gfz5mk
- Simon Kallweit, Petrik Clarberg, Craig Kolb, Tom'as Davidović, Kai-Hwa Yao, Theresa Foley, Yong He, Lifan Wu, Lucy Chen, Tomas Akenine-Möller, Chris Wyman, Cyril Crassin, and Nir Benty. 2022. The Falcor Rendering Framework. <https://github.com/NVIDIAGameWorks/Falcor>
- Markus Kettunen, Marco Manzi, Miika Aittala, Jaakko Lehtinen, Frédo Durand, and Matthias Zwicker. 2015. Gradient-Domain Path Tracing. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 34, 4 (July 2015), 123. doi:10/gfzrhn
- Juhyeon Kim, Craig Benko, Magnus Wrenninge, Ryusuke Villemin, Zeb Barber, Wojciech Jarosz, and Adithya Pediredla. 2025. A Monte Carlo Rendering Framework for Simulating Optical Heterodyne Detection. *ACM Transactions on Graphics (TOG)* 44, 4 (2025), 1–19.
- Juhyeon Kim, Wojciech Jarosz, Ioannis Gkioulekas, and Adithya Pediredla. 2023. Doppler Time-of-Flight Rendering. *ACM Transactions on Graphics (TOG)* 42, 6 (2023), 1–18.
- Juhyeon Kim, Wojciech Jarosz, and Adithya Pediredla. 2026. ToF ReSTIR: Time-of-Flight Rendering with Spatio-temporal Reservoir Resampling. *ACM Transactions on Graphics (TOG)* 45, 4 (2026), 1–18.
- You Li and Javier Ibanez-Guzman. 2020. Lidar for autonomous driving: The principles, challenges, and trends for automotive lidar and perception systems. *IEEE Signal Processing Magazine* 37, 4 (2020), 50–61.
- Daqi Lin, Markus Kettunen, Benedikt Bitterli, Jacopo Pantaleoni, Cem Yuksel, and Chris Wyman. 2022. Generalized Resampled Importance Sampling: Foundations of ReSTIR. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 41, 4 (July 2022), 75:1–75:23. doi:10/ggjn7b
- Yang Liu, Shaojie Jiao, and Wojciech Jarosz. 2022. Temporally Sliced Photon Primitives for Time-of-Flight Rendering. *Computer Graphics Forum (Proceedings of the Eurographics Symposium on Rendering)* 41, 4 (2022). doi:10.1111/cgf.14584
- Marco Manzi, Markus Kettunen, Frédo Durand, Matthias Zwicker, and Jaakko Lehtinen. 2016. Temporal Gradient-Domain Path Tracing. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 35, 6 (Dec. 2016), 246:1–246:9. doi:10/f9cpw
- Julio Marco, Wojciech Jarosz, Diego Gutierrez, and Adrian Jarabo. 2017. Transient Photon Beams. In *Congreso Espanol de Informatica Grafica*. Eurographics Association. doi:10/gfznc
- Esdras Medeiros, Harish Doraiswamy, Matthew Berger, and Claudio T Silva. 2014. Using physically based rendering to benchmark structured light scanners. In *Computer Graphics Forum*, Vol. 33. Wiley Online Library, 71–80.
- Shree K Nayar, Gurunandan Krishnan, Michael D Grossberg, and Ramesh Raskar. 2006. Fast Separation of Direct and Global Components of a Scene Using High Frequency Illumination. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 25, 3 (2006), 935–944. doi:10/fw3cpp
- Adithya Pediredla, Ashok Veeraraghavan, and Ioannis Gkioulekas. 2019. Ellipsoidal Path Connections for Time-Gated Rendering. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 38, 4 (July 2019), 38:1–38:12. doi:10/gf5jbm
- Joaquim Salvi, Jordi Pages, and Joan Batlle. 2004. Pattern codification strategies in structured light systems. *Pattern recognition* 37, 4 (2004), 827–849.
- Chenbo Shi, Zheng Qin, Xiaowei Hu, Changsheng Zhu, Yuanzheng Mo, Zelong Li, Shaojia Yan, Yue Yu, Xiangteng Zang, and Chun Zhang. 2024. Phase error reduction for a structured-light 3D system based on a texture-modulated reprojection method. *Sensors* 24, 7 (2024), 2075.
- Kartic Subr, Derek Nowrouzezahrai, Wojciech Jarosz, Jan Kautz, and Kenny Mitchell. 2014. Error Analysis of Estimators That Use Combinations of Stochastic Sampling Strategies for Direct Illumination. *Computer Graphics Forum (Proceedings of the Eurographics Symposium on Rendering)* 33, 4 (June 2014), 93–102. doi:10/f6fgw4
- Eric Veach. 1997. *Robust Monte Carlo Methods for Light Transport Simulation*. Ph.D. Dissertation. Stanford University.
- Lifan Wu, Guangyan Cai, Ravi Ramamoorthi, and Shuang Zhao. 2021. Differentiable Time-Gated Rendering. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 40, 6 (Dec. 2021). doi:10/h2c9
- Bing Xu, Tzu-Mao Li, Iliyan Georgiev, Trevor Hedstrom, and Ravi Ramamoorthi. 2024. Residual path integrals for re-rendering. In *Computer Graphics Forum*, Vol. 43. Wiley Online Library, e15152.
- Zihan Yu, Cheng Zhang, Derek Nowrouzezahrai, Zhao Dong, and Shuang Zhao. 2022. Efficient Differentiation of Pixel Reconstruction Filters for Path-Space Differentiable Rendering. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 41, 6 (Nov. 2022), 191:1–191:16. doi:10/grqv42
- Pietro Zanuttigh, Giulio Marin, Carlo Dal Mutto, Fabio Dominio, Ludovico Minto, Guido Maria Cortelazzo, et al. 2016. Time-of-flight and structured light depth cameras. *Technology and Applications* 978, 3 (2016).
- Cheng Zhang, Zhao Dong, Michael Doggett, and Shuang Zhao. 2021. Antithetic Sampling for Monte Carlo Differentiable Rendering. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 40, 4 (July 2021), 77:1–77:12. doi:10/gmmx7
- Chao Zuo, Shijie Feng, Lei Huang, Tianyang Tao, Wei Yin, and Qian Chen. 2018. Phase shifting algorithms for fringe projection profilometry: A review. *Optics and lasers in engineering* 109 (2018), 23–59.




Fig. 10. Equal-time comparison of modulated indirect illumination rendered using naive and antithetic sampling across multiple scenes and wavelengths. Per-image MSE values are reported above each result. Antithetic sampling substantially reduces noise at short wavelengths, while its advantage decreases at longer wavelengths as the required geometric deformation increases. Complete MSE curves are shown in Fig. 4.




Fig. 11. Fringe projection profilometry for automated optical inspection. (A) Rendered RGB image of a circuit board. (B–C) Depth reconstructions obtained using naive and antithetic sampling. (D) Ground-truth depth. Depth is recovered from four phase-shifted sinusoidal patterns through phase decoding, spatial unwrapping, and projector–camera triangulation. All depth images use the same colormap and range. Antithetic sampling reduces indirect-illumination noise, improving phase stability and reconstruction accuracy. Residual deviations from ground truth arise from global illumination and texture-edge artifacts, which also occur in real-world measurements [Deng et al. 2020; Hu et al. 2023; Shi et al. 2024].




Fig. 12. (A) Direct–global illumination decomposition using high-frequency checkerboard patterns [Nayar et al. 2006] for diffuse, translucent, and specular objects. The reconstructed direct and indirect illumination components are compared with the ground truth. (B) Error maps for naive and antithetic sampling. By reducing noise in the rendered checkerboard measurements, antithetic sampling produces more accurate decompositions across different material types.




Fig. 13. Comparison of antithetic mappings for 50 arbitrary binary patterns of length 1024. (A) Rendering MSE versus Earth Mover’s Distance (EMD) for naive sampling and antithetic sampling with scale-based and optimal-transport (OT) mappings. OT performs best for low-EMD patterns, while scale-based mapping becomes preferable for patterns with higher EMD; its performance depends instead on the mapping Jacobian. (B) Rendered results and corresponding OT mappings for representative low- and high-EMD patterns, illustrating how larger transport distances reduce path similarity and increase rendering error.