Non-line of Sight Transient Rendering

Diego Royo[1], Jorge Garcia[1], Adolfo Muñoz[1], Adrián Jarabo[2]

[1]Affiliation: Graphics and Imaging Lab (GILab)
Instituto de Investigación en Ingeniería de Aragón (I3A)
Universidad de Zaragoza, Mariano Esquillor s/n, 50018, Zaragoza, Spain.
Tel. +34-976762707, e-mail: *droyo@unizar.es*
[2]Affiliation: Meta Reality Labs, Zürich, Switzerland

Abstract

Non-line of sight imaging reconstructs hidden geometry using time-resolved measurements of indirect diffuse reflections emitted by a laser. To help develop new applications, we simulate light transport in these occluded scenes and introduce a set of simple, yet effective subpath sampling techniques for efficiency, with an open-source implementation.

Introduction

Transient imaging aims to capture and analyze how light propagates through a scene with ultra-high temporal resolution, reaching the order of picoseconds. We focus on non-line-of-sight (NLOS) imaging, which attempts to reconstruct fully or partially occluded scenes by measuring the back-scattering of the hidden objects in a secondary visible relay wall (Figure 1). This has a wide range of applications in autonomous driving, medical imaging, space exploration, and others. These scenarios usually require expensive hardware, including pulsed illumination and ultra-fast cameras, which are difficult to calibrate and operate. Transient light transport simulation thus emerges as an alternative tool for developing, prototyping, and testing such systems, before building them.

In our work, we build upon the transient path integral formulation proposed by Jarabo et al. [3]. In a NLOS setup, the geometry of interest and light sources are hidden behind multiple bounces in the scene, so light paths generated with a conventional algorithm are sub-optimal for this purpose. Therefore, in this work, we develop three simple yet effective sampling techniques for generating eye, light and hidden geometry subpaths that produce intermediate *virtual* irradiance meters and light sources. These new techniques improve the convergence of NLOS transient path tracing (see Figure 2) with very little computational overhead.

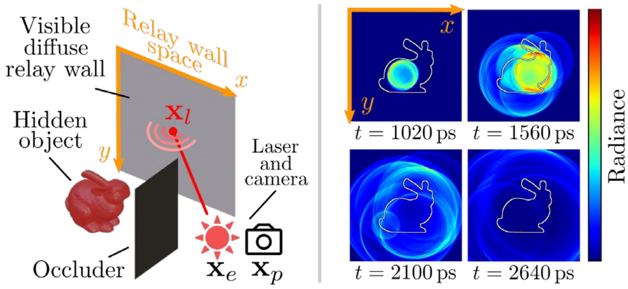

Figure 1: Left: A laser pulse is emitted towards a visible diffuse wall, illuminating the occluded objects. Its reflection hits the relay wall and is captured by an ultra-fast time-resolved camera. Right: rendered simulation.

We implement our methods in the Mitsuba 2 rendering system [5], using our extension for transient rendering, which provides many advantages over previous transient path tracing systems [2, 3], including support for CPU and GPU parallelization, simulation of the polarization of light, and the possibility of extending its differentiable rendering capabilities to the NLOS domain. Although there exist other NLOS capture simulation methods [1], up to our knowledge our implementation is the first comprehensive transient extension to Mitsuba 2 available for public[1] use.

Our approach

Our sampling strategies extend the path contribution function for transient path tracing [3]. Consider a path $\bar{\mathbf{x}} = \mathbf{x}_0 \dots \mathbf{x}_k \in \Omega$ defined as a sequence of vertices from the light source $\mathbf{x}_0$ to the camera sensor $\mathbf{x}_k$. The light measured for that path is proportional to:

$$f(\bar{\mathbf{x}}) \propto L_e(\mathbf{x}_0 \to \mathbf{x}_1)\, W_e(\mathbf{x}_{k-1} \to \mathbf{x}_k), \quad (1)$$

where L_e is the light emitted from the source $\mathbf{x}_0$ towards $\mathbf{x}_1$ and the sensor importance W_e models spatial and angular sensitivity in the same way. Also, the probability of each vertex corresponding to a forward path tracer can be expressed as:

$$p(\bar{\mathbf{x}}) = p(\mathbf{x}_k) \prod_{i=0}^{k-1} p(\mathbf{x}_i | \mathbf{x}_{i+1}). \quad (2)$$

[1] https://github.com/diegoroyo/mitsuba2-transient-nlos

Relay wall space

Consider a time-resolved image $I(x, y, t)$ where t denotes the temporal domain. The key difference between conventional line-of-sight (LOS) images and their NLOS counterparts lies in the topology of (x, y). While LOS images aim to reconstruct a continuous signal, NLOS captures use sparse points in the visible relay wall, with sensors that only measure light coming from a differential solid angle. We define a mapping from a parameterized 2D relay wall coordinates (x_s, y_s) to its corresponding 3D point as $\mathcal{M}_s(x_s, y_s) \to \mathbf{x}_s$. Thus, W_e becomes:

$$W_e(\mathbf{x}_{k-1} \to \mathbf{x}_k) = \delta(\mathbf{x}_p - \mathbf{x}_k) \, \delta(\omega_s - \omega_{k-1}), \tag{3}$$

with $\delta(\cdot)$ the Dirac delta, $\mathbf{x}_p$ the position of the sensor directed at $\omega_s = \text{norm}(\mathcal{M}_s(x_s, y_s) - \mathbf{x}_p)$ and $\omega_{k-1} = \text{norm}(\mathbf{x}_{k-1} - \mathbf{x}_k)$ normalized directions. The distribution of these vertices for each path, i.e. $p(\mathbf{x}_{k-1}|\mathbf{x}_k)$ is uniform and regular along the relay wall parameterized 2D space.

Laser sampling

In a typical NLOS capture setup, light is emitted from a laser in a differential solid angle towards a single point in the visible relay wall, defined as $\mathcal{M}_l(x_l, y_l) = \mathbf{x}_l$ with $\mathcal{M}_l$ similar to $\mathcal{M}_s$. That is,

$$L_e(\mathbf{x}_0 \to \mathbf{x}_1) = \delta(\mathbf{x}_e - \mathbf{x}_0) \, \delta(\omega_l - \omega_0), \tag{4}$$

with $\mathbf{x}_e$ the laser position, $\omega_0 = \text{norm}(\mathbf{x}_0 - \mathbf{x}_1)$ and $\omega_l = \text{norm}(\mathbf{x}_e - \mathcal{M}_l(x_l, y_l))$. Our laser sampling technique is similar to next event estimation, targeting $\mathbf{x}_l$ instead of $\mathbf{x}_e$. We can analytically compute the light contribution $\hat{L}(\mathbf{x}_0 \to \mathbf{x})$ to any point $\mathbf{x}$ through $\mathbf{x}_l$ with probability $p(\mathbf{x}_1|\mathbf{x}_2) = \delta(\mathbf{x}_l - \mathbf{x}_1)$. This method ignores longer paths (e.g. interreflections) and should be combined with classical sampling strategies in multiple importance sampling.

Hidden geometry sampling

The back-scattering of the hidden geometry can only be found through an indirect bounce on the relay wall. In a forward path tracer, the generated vertices follow a distribution $p(\mathbf{x}_{j-1}|\mathbf{x}_j)$ which targets the reflectance distribution or the light source location, both suboptimal in NLOS scenes.

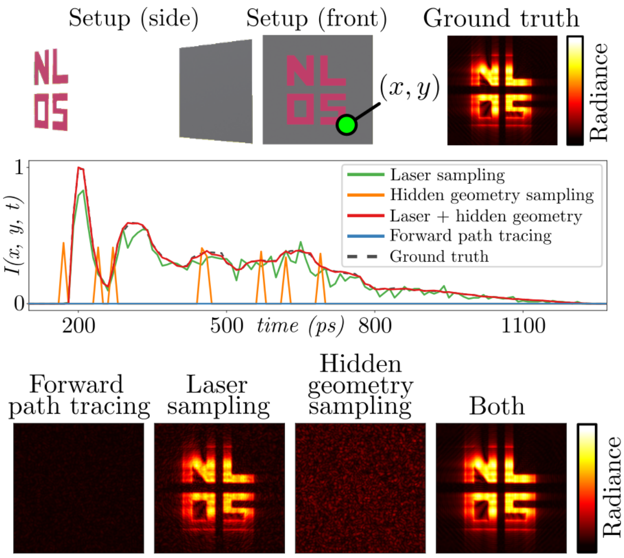

Figure 2: Ablation test for our sampling techniques, 200,000 samples per pixel with 256 x 256 resolution. Top: NLOS letters scene. Middle: time-resolved radiance for the (x, y) pixel. Without laser sampling, we use a light projector with a small solid angle to illuminate a tiny area in the relay wall. The ground truth has been obtained using several orders of magnitude more samples per pixel. Bottom: phasor fields-based reconstruction of hidden geometry [4].

For a path, we generate a random vertex $\mathbf{x}$ on the surface of the hidden geometry with uniform probability

$$p(\mathbf{x}_{k-2}|\mathbf{x}_{k-1}) = 1/A, \tag{5}$$

where A is the surface area of the hidden geometry.

ACKNOWLEDGEMENTS

This work has received funding from the European Research Council (ERC) under the European Union's Horizon 2020 research and innovation programme (project CHAMELEON, Grant No 682080), the EU MSCA-ITN programme (project PRIME, Grant No 956585) and Spain's Agencia Estatal de Investigación (project PID2019-105004GB-I00). Additionally, Diego Royo was supported by a Gobierno de Aragón (2022-2026) predoctoral grant.

REFERENCES

[1] ISERINGHAUSEN, I. and HULLIN, M.B., 2020. Non-line-of-sight reconstruction using efficient transient rendering. *ACM Transactions on Graphics*. 39(1), 1–14.

[2] JARABO, A. and ARELLANO, V, 2018.. Bidirectional rendering of vector light transport. *Computer Graphics Forum*. 37(6), 96-105.

[3] JARABO, A., et al., 2014. A framework for transient rendering. *ACM Transactions on Graphics*. 33(6), 1-10.

[4] LIU, X., et al., 2019. Non-line-of-sight imaging using phasor-field virtual wave optics. *Nature*. 572(7771), 620-623.

[5] NIMIER-DAVID, M., et al., 2019. Mitsuba 2: A retargetable forward and inverse renderer. *ACM Transactions on Graphics*. 38(6), 1-17.