

Ghost ReSTIR: Volumetric Resampling with Ghost Vertices in Null-Scattering Space

Song Zhang
University of Utah
Salt Lake City, USA
song.zhang@utah.edu

Markus Kettunen
NVIDIA Research
Helsinki, Finland
mkettunen@nvidia.com

Daqi Lin
NVIDIA Research
Redmond, USA
daqil@nvidia.com

Cem Yuksel
University of Utah
Salt Lake City, USA
cem@cemyuksel.com

Pengpei Hong
University of Utah
Salt Lake City, USA
pengpei.hong@utah.edu

Chris Wyman
NVIDIA Research
Redmond, USA
cwyman@nvidia.com




Figure 1: Ghost ReSTIR improves both quality and performance over null-scattering path tracing [Galtier et al. 2013; Miller et al. 2019] and prior volumetric ReSTIR [Lin et al. 2021]. (Left) **EXPLOSION**, a dense emissive volume under 9-bounce multiple scattering and dynamic camera—prior volumetric ReSTIR shows severe noise in dense emissive regions from its inconsistent $\hat{p}$, which our consistent formulation substantially reduces. (Right) **BISTRO**, a smoke plume lit by many colored area lights and surfaces in single scattering—ours handles volume–surface interaction more cleanly and runs 1.4× faster thanks to null-scattering tracking that avoids the costly voxel traversal of regular tracking.

Abstract

Spatiotemporal reservoir resampling (ReSTIR) has been extended to volumetric media, achieving significant quality improvements under arbitrary dynamic illumination. However, it still suffers from inconsistent transmittance approximations across resampling stages, and performance scales poorly with volume resolution.

We present *Ghost ReSTIR*, a new volumetric ReSTIR method formulated in null-scattering primary sample space. Our target



This work is licensed under a Creative Commons Attribution 4.0 International License.
SIGGRAPH Posters '26, Los Angeles, CA, USA
© 2026 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2548-7/26/07
<https://doi.org/10.1145/3799825.3818719>

function contains no intermediate transmittance, enabling consistent evaluation across all resampling stages. We introduce ghost vertices—auxiliary infinite tails appended to each path segment—that provide a well-defined bijection for reconnection shifts regardless of null-vertex count or density changes, removing the need to store historical volume data for dynamic scenes. Our formulation further enables MIS between next-event estimation and phase/BSDF sampled paths. Across dense volumetric scenes, Ghost ReSTIR improves resampling quality and runs up to 30% faster than prior volumetric ReSTIR, while requiring no mipmap tuning for transmittance approximation.

CCS Concepts

• Computing methodologies → Rendering.

Keywords

real-time path tracing, resampled importance sampling, ReSTIR, null-scattering, volume rendering

ACM Reference Format:

Song Zhang, Daqi Lin, Pengpei Hong, Markus Kettunen, Cem Yuksel, and Chris Wyman. 2026. Ghost ReSTIR: Volumetric Resampling with Ghost Vertices in Null-Scattering Space. In *Special Interest Group on Computer Graphics and Interactive Techniques Conference Posters (SIGGRAPH Posters '26)*, July 19–23, 2026, Los Angeles, CA, USA. ACM, New York, NY, USA, 3 pages. <https://doi.org/10.1145/3799825.3818719>

1 Introduction

Participating media are among the most expensive components of physically based rendering, especially in large, dense, dynamic volumes. ReSTIR methods greatly improve sampling efficiency for surfaces [Bitterli et al. 2020; Lin et al. 2022], and Lin et al. [2021] extended this to volumetric transport. However, prior volumetric ReSTIR has three key limitations:

- (1) **Inconsistent target function.** Lin et al. [2021] include transmittance T in the target $\hat{p}$, but evaluating T in heterogeneous media is expensive, so they use progressively refined approximations across stages: piecewise-constant T^* for candidate generation, ray-marched $\tilde{T}$ during reuse, and unbiased T only at shading. In high-frequency density regions, T^* and $\tilde{T}$ deviate substantially from T , so resampling weights poorly reflect the true path contribution, causing excess variance in dense media.
- (2) **Performance scales with grid resolution and density changes.** Built on regular tracking, whose cost depends on traversing the voxel grid, performance degrades as resolution increases and density decreases (longer paths).
- (3) **Limited shift mappings.** It predates GRIS [Lin et al. 2022], using direction-and-distance copy shifts rather than hybrid replay-plus-reconnection shifts, and cannot combine NEE and phase/BSDF paths via MIS.

We address all three by reformulating volumetric ReSTIR in *null-scattering primary sample space* (PSS). Null scattering [Galtier et al. 2013; Miller et al. 2019] samples free-flight distances from a spatially simple majorant rather than traversing the voxel grid, improving scalability. In PSS, all transmittance terms cancel between the integrand and sampling PDF, yielding a transmittance-free $\hat{p}$ consistent

across all stages. We further introduce *ghost vertices*—auxiliary infinite tails extending the path space—that enable well-defined reconnection shifts under variable null-vertex counts and dynamic media, and support MIS between NEE and phase/BSDF paths.

2 Ghost ReSTIR

2.1 Ghost Vertex Formulation in Primary Sample Space

The null-scattering framework [Miller et al. 2019] expresses the pixel measurement as a path integral $I = \int_{\mathcal{P}} f(\bar{\mathbf{x}}) d\bar{\mathbf{x}}$, where $f(\bar{\mathbf{x}})$ is the path’s measurement contribution. Each segment between two real vertices (surface or volume scattering events) may contain a variable number of null vertices. This variable dimensionality prevents direct reconnection shifts, since base and offset segments generally have different null-vertex counts.

We resolve this with a *ghost volume* $\mathcal{V}_g$ spanning all of $\mathbb{R}^3$ with a positive majorant $\bar{\mu}_g(\mathbf{x}) > 0$ everywhere. After each real vertex $\mathbf{x}_{r_i}$ (excluding the camera), we append an infinite sequence of ghost vertices along the incoming ray direction. The augmented path $\hat{\mathbf{x}}$ is the original path $\bar{\mathbf{x}}$ plus one infinite ghost tail per real vertex, and its integrand factorizes as:

$$f(\hat{\mathbf{x}}) = f(\bar{\mathbf{x}}) \cdot \prod_{i=1}^r \prod_{j=1}^{\infty} \bar{T}_g(\mathbf{x}_{r_{i-1}}^i, \mathbf{x}_{r_i}^i) \bar{\mu}_g(\mathbf{x}_{r_i}^i), \quad (1)$$

where each factor $\bar{T}_g \cdot \bar{\mu}_g$ is exactly the free-flight PDF under $\bar{\mu}_g$. Since $\bar{\mu}_g > 0$ everywhere and the ray extends to infinity, integrating out each ghost vertex yields unity, so $\int f(\hat{\mathbf{x}}) d\hat{\mathbf{x}} = \int f(\bar{\mathbf{x}}) d\bar{\mathbf{x}}$: ghost vertices leave the image unchanged but add the degrees of freedom needed for shift mappings.

In PSS, the integrand becomes the ratio $F(\hat{\mathbf{u}}) = f(\hat{\mathbf{x}})/p(\hat{\mathbf{x}})$, where all transmittance, geometry, and ghost-tail terms cancel, leaving only scattering-to-sampling ratios at real vertices and emitted radiance. The target $\hat{p} = \|F\|$ thus contains no transmittance and is evaluated consistently across all stages without approximation.

When generating initial candidates, ghost vertices need not be sampled explicitly: they cancel in F and sampling proceeds as standard delta tracking. They matter only during shift mapping, and the formulation naturally supports technique MIS between NEE and phase/BSDF paths.

2.2 Hybrid Shift with Ghost Vertices

The hybrid shift combines random replay for early bounces with reconnection at a later vertex. *Random replay* copies all random numbers as in ReSTIR PT [Lin et al. 2022], with the ghost tails absorbing any difference in null-vertex counts (Jacobian = 1). *Reconnection* is more involved.

Consider one reconnection segment. In the base path’s distance sequence $\mathbf{d} = [u_0, u_1, \dots]$, entries $u_0, \dots, u_{n-1}$ sample the n null vertices before the shared reconnection vertex $\mathbf{x}_r$, u_n samples the real collision at $\mathbf{x}_r$, and $u_{n+1}, \dots$ sample the infinite ghost tail. We remove u_n and *reparameterize* u_{n+1} so the remaining subsequence replays continuously from any preceding vertex:

$$\tilde{u}_{n+1} = \bar{T}(\mathbf{x}_{r-1}, \mathbf{x}_r) (u_{n+1} - 1) + 1, \quad (2)$$




Figure 2: Highly anisotropic static PLUME volume ($g = -0.95$) lit by a high-contrast SHANGHAI environment map, with near equal-time comparison. *Baseline* cannot use MIS between NEE and phase-sampled paths, leading to severe correlation artifacts. Our formulation enables technique MIS, substantially reducing both noise and correlation compared to *Null-Scattering* and *Baseline*. The scene is captured under camera motion.

accounting for the optical depth between the last null vertex $\mathbf{x}_{r-1}$ and $\mathbf{x}_r$. The reparameterized subsequence $\tilde{\mathbf{d}}$ —the original $\mathbf{d}$ with u_n removed and u_{n+1} replaced by $\tilde{u}_{n+1}$ —is copied to the offset path and replayed from its real vertex $\mathbf{x}'_{r-}$. Since the offset segment differs in length, replay yields a different null count n' , set by when accumulated distance first reaches $\mathbf{x}_r$; collisions that would give a real event before $\mathbf{x}_r$ are forced to null (as in ratio tracking). The inverse reparameterization then converts back to the offset’s original sequence.

The resulting reconnection Jacobian is:

$$\left| \frac{\partial \mathbf{u}'}{\partial \mathbf{u}} \right| = \left| \frac{\partial \mathbf{e}'}{\partial \mathbf{e}} \right| \cdot \frac{\tilde{\mu}_t(\mathbf{x}'_r) p(\omega' | \mathbf{x}'_{r-}) G(\mathbf{x}'_{r-}, \mathbf{x}_r)}{\tilde{\mu}_t(\mathbf{x}_r) p(\omega | \mathbf{x}_{r-}) G(\mathbf{x}_{r-}, \mathbf{x}_r)}, \quad (3)$$

where $|\partial \mathbf{e}' / \partial \mathbf{e}|$ accounts for event-type remapping between null and ghost vertices, and the remaining terms involve only geometry, direction PDFs, and majorant ratios. Crucially, *no transmittance appears in the Jacobian*—the reparameterization is designed for this cancellation. The reservoir stores only three extra values per reconnection segment (a local seed, the reparameterized optical depth, and a ghost seed), keeping storage constant regardless of null count.

3 Results

We measure performance on an NVIDIA RTX 4090 at 1920×1080 , reporting mean absolute percentage error (MAPE), against two baselines: *Null-Scattering* [Galtier et al. 2013; Miller et al. 2019] (null-scattering path tracing with delta and ratio tracking) and *Baseline* [Lin et al. 2021] (prior volumetric ReSTIR, ported to Falcor 8). All methods use GVDB [Hoetzlein 2016] for VDB access.

We evaluate several volumetric scenes—EXPLOSION, BISTRO SMOKE (volume + surfaces), BUNNY CLOUD, and EMERALD SQUARE FOG (animated volume + surfaces)—under single and multiple scattering (see Figure 1 and supplementary video). Ours consistently achieves the lowest MAPE with visibly less noise, most notably where density or emission varies rapidly (EXPLOSION). It runs 10% faster than *Baseline* on larger volumes (BUNNY CLOUD) and 26% faster on volume-surface scenes (BISTRO SMOKE), with a slight slowdown on the

smaller PLUME grid—thanks to null-scattering tracking, whose cost is independent of grid resolution, plus early reconnection that shortens shift evaluation. Our PSS formulation also supports technique MIS between NEE and phase/BSDF-sampled paths, substantially reducing the noise and correlation artifacts *Baseline* shows in highly anisotropic media ($g = -0.95$), where it can only rely on NEE (Figure 2).

References

- 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, Article 148 (July 2020), 148:1–148:17 pages. doi:10.1145/3386569.3392481
- Mathieu Galtier, Stéphane Blanco, Cyril Caliot, Christophe Coustet, Jérémie Dauchet, Mouna El Hafi, Vincent Eymet, Richard Fournier, Jacques Gautrais, Anaïs Khuong, Benjamin Piau, and Guillaume Terrée. 2013. Integral Formulation of Null-Collision Monte Carlo Algorithms. *Journal of Quantitative Spectroscopy and Radiative Transfer* 125 (2013), 57–68. doi:10.1016/j.jqsrt.2013.04.001
- Rama Karl Hoetzlein. 2016. GVDB: Raytracing Sparse Voxel Database Structures on the GPU. In *Proceedings of High Performance Graphics*. 109–117.
- 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, Article 75 (July 2022), 75:1–75:23 pages. doi:10.1145/3528223.3530158 Joint first authors: Daqi Lin and Markus Kettunen.
- Daqi Lin, Chris Wyman, and Cem Yuksel. 2021. Fast Volume Rendering with Spatiotemporal Reservoir Resampling. *ACM Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 40, 6, Article 279 (Dec. 2021), 279:1–279:18 pages. doi:10.1145/3478513.3480499
- Bailey Miller, Iliyan Georgiev, and Wojciech Jarosz. 2019. A Null-Scattering Path Integral Formulation of Light Transport. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 38, 4, Article 44 (July 2019), 44:1–44:13 pages. doi:10.1145/3306346.3323025