Sample Space Partitioning and Spatiotemporal Resampling for Specular Manifold Sampling

PENGPEI HONG, University of Utah, USA MENG DUAN, Nankai University, China BEIBEI WANG, Nanjing University, China CEM YUKSEL, University of Utah, USA TIZIAN ZELTNER, NVIDIA, Switzerland DAQI LIN, NVIDIA, USA

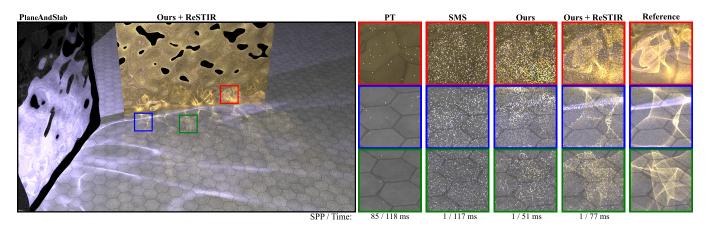


Fig. 1. A comparison between our method, our method with ReSTIR, path tracing (PT), and specular manifold sampling (SMS) in a scene containing a two-sided refractive slab and a metallic reflective plane, both with normal variation to create caustic patterns. The scene is lit by two different-colored area lights. With only 1 sample per pixel, our method significantly reduces the variance and the frame time compared to SMS. Combining our method with ReSTIR's spatiotemporal resampling produces a more converged rendering in real time. Our sample space partitioning uses 32×32 tiles and 3×3 trust regions.

Caustics rendering remains a long-standing challenge in Monte Carlo rendering because high-energy specular paths occupy only a small region of path space, making them difficult to sample effectively. Recent work such as Specular Manifold Sampling (SMS) [Zeltner et al. 2020] can stochastically sample these specular paths and estimate their unbiased weights using Bernoulli trials. However, applying SMS in interactive rendering is non-trivial because it is slow and delivers noisy images given a very limited time budget.

In this work, we extend SMS for high-quality caustic rendering in interactive settings using *sample space partitioning*. Our insight is that Newton iterations, the main performance bottleneck of SMS, can be restricted to the vicinity of the seed path, which can dramatically improve the performance. We achieve this with tile-based sample space partitioning, which bounds the manifold walk region and allows building a per-frame prior distribution

Authors' Contact Information: Pengpei Hong, University of Utah, Salt Lake City, USA, u1472289@utah.edu; Meng Duan, Nankai University, Tianjin, China, mengduan824@ gmail.com; Beibei Wang, Nanjing University, Suzhou, China, beibei.wang@nju.edu.cn; Cem Yuksel, University of Utah, Salt Lake City, USA, cem@cemyuksel.com; Tizian Zeltner, NVIDIA, Zurich, Switzerland, tizian.zeltner@gmail.com; Daqi Lin, NVIDIA, Redmond, USA, daqil@nvidia.com.



This work is licensed under a Creative Commons Attribution 4.0 International License. SA Conference Papers '25, Hong Kong, Hong Kong
© 2025 Copyright held by the owner/author(s).
ACM ISBN 979-8-4007-2137-3/25/12
https://doi.org/10.1145/3757377.3763927

that concentrates initial guesses around solutions. This reduces the cost of SMS and improves its sampling quality. Applying spatiotemporal reuse (ReSTIR) further amortizes the sample generation cost, greatly increasing the effective sample count. As a result, we achieve significant variance reduction compared to SMS in interactive rendering scenarios.

CCS Concepts: • Computing methodologies → Ray tracing.

Additional Key Words and Phrases: path tracing, manifold walk, Newton iterations, Bernoulli trials, resampling

ACM Reference Format:

Pengpei Hong, Meng Duan, Beibei Wang, Cem Yuksel, Tizian Zeltner, and Daqi Lin. 2025. Sample Space Partitioning and Spatiotemporal Resampling for Specular Manifold Sampling. In SIGGRAPH Asia 2025 Conference Papers (SA Conference Papers '25), December 15–18, 2025, Hong Kong, Hong Kong. ACM, New York, NY, USA, 10 pages. https://doi.org/10.1145/3757377.3763927

1 Introduction

Caustics, light patterns produced by specular materials, have been a challenge to render efficiently. Recent Specular Manifold Sampling (SMS) [Zeltner et al. 2020] work provides a principled and unbiased approach to directly sample caustic paths, overcoming several limitations of prior techniques. For instance, bidirectional path tracing [Veach and Guibas 1994, 1997] struggles with specular-diffuse-specular (SDS) paths, while photon mapping methods [Hachisuka

and Jensen 2009; Jensen 1996; Kern et al. 2024] introduce blurring bias due to density estimation.

In this paper, we extend SMS for interactive rendering of caustics, simultaneously improving its speed and quality. While SMS can provide high-quality caustics for offline rendering, it incurs a substantial amount of computation cost due to repeated Newton iterations and the ray tracing it requires for identifying caustic paths and estimating their sampling probabilities within a 2D primary sample space. With limited sample budgets in interactive rendering, SMS can exhibit excessive noise.

To minimize the cost of SMS, we introduce *sample space partitioning*, which divides the primary sample space into *tiles* to bound the area considered during Newton iterations. Thereby, we significantly reduce the number of Newton iterations (and the corresponding ray tracing) by not only bounding the solution space but also requiring fewer samples to estimate the sampling probability of found paths.

However, using tiles increases the variance when the convergence basins—regions of initial guesses that converge to the same caustic path—are large. We rectify this by constructing a per-frame prior distribution via presampling the SMS solutions from a subset of nearby pixels to identify convergence basin regions. Due to spatial continuity, we can identify tiles where solutions are likely to exist for a block of pixels. This prior enables importance sampling, allowing us to start at the vicinity of solutions and improve the success rate of finding a solution, thereby both reducing the variance and improving the performance compared to SMS.

Finally, we apply spatiotemporal resampling (ReSTIR) [Bitterli et al. 2020] to aggregate samples from different prior distributions across frames to improve the final sample per frame towards a target distribution that accounts for the energy of the light sources and the attenuation effect of specular paths. ReSTIR helps amortizing the cost of sample generation and inverse probability estimation over frames as it only requires manifold shifts for reusing samples.

We evaluate our method in various scenes, including setups with multiple reflective and refractive objects and light sources (Figure 1), and observe substantial improvements over SMS, achieving faster render times and significantly lower noise levels.

In summary, our main contributions are:

- *Tile-based sample space partitioning* that constrains the Newton solver to accelerate both specular path finding and inverse probability estimation (Section 3).
- Per-frame prior distributions over tiles to importance sample seed paths around solutions to achieve higher sampling efficiency (Section 4).
- Integration of spatiotemporal resampling that amortizes sample generation cost and enables rendering caustics at higher effective sample counts (Section 5).

2 Related Work

Caustics, especially ones formed by indirect specular-diffuse-specular (SDS) transport, are difficult to render with common path-space sampling techniques due to their high energy concentration in a small region. We first review prior works on caustic rendering and discuss their suitability for interactive rendering.

Bidirectional methods. Bidirectional path tracing [Veach and Guibas 1994] traces paths starting from light sources and can directly splat the contribution from caustic paths to the image. However, it cannot handle SDS paths. This limitation is overcome by photon mapping [Jensen 1996], which has been widely used to render caustics in production renderers [Burley et al. 2018; D5Render 2023]. Photon mapping applies a blurring kernel to share incident radiance of photons between nearby shading points. A consistent estimator can be formed by progressively shrinking the kernel radius [Georgiev et al. 2012; Hachisuka and Jensen 2009]. A limitation of bidirectional methods is that light subpaths are not importance sampled based on the camera's viewpoint, which can reduce efficiency in complex scenes as many photons may contribute little to the final image.

Path guiding, MLT, and resampling. A different category of methods is based on reusing information found from previous samples. Path guiding methods use training samples to build local distributions of incident radiance which are represented as histograms [Müller et al. 2017], Gaussian/vMF mixtures [Ruppert et al. 2020; Vorba et al. 2014], or a neural network [Dong et al. 2023; Müller et al. 2018]. These methods can learn the incoming radiance distribution very well and are widely adopted in offline rendering. Problems of path guiding methods for interactive caustic rendering include insufficient training samples, limited angular resolution, and the expensive data structures they require. Metropolis Light Transport (MLT) [Veach and Guibas 1997] and its variants [Hachisuka et al. 2014; Kaplanyan et al. 2014; Kelemen et al. 2002] can find high contribution paths by exploring path space using MCMC mutations, but have temporal instability that hinders their practical use [Fascione et al. 2018]. ReSTIR [Bitterli et al. 2020; Lin et al. 2022] resamples candidates spatiotemporally to approximate a target distribution in real-time. However, it can spread high-energy fireflies of outlier candidates into correlation blobs [Sawhney et al. 2024], which is particularly problematic if the chance of finding caustics is too low.

Manifold walks and analytical methods. Originally developed as an MCMC mutation of valid specular paths, manifold walks [Jakob and Marschner 2012] allow solving for an admissible specular path from an initially invalid path using an iterative Newton solver scheme. Manifold next-event estimation (MNEE) [Hanika et al. 2015] initializes manifold walks using a straight line connecting a shading point with a light source, enabling a unidirectional path tracer to handle both direct and indirect (SDS) caustics, albeit finding at most one solution. Specular manifold sampling (SMS) [Zeltner et al. 2020] generalizes MNEE by randomizing the choice of seed paths, ensuring all admissible paths can be sampled. However, SMS is not directly compatible with interactive rendering: The unbiased variant of SMS requires an unbounded number of Bernoulli trials to determine the (inverse) probability of found solutions, while the biased variant requires many samples to avoid energy loss.

A key problem of SMS is the high variance caused by its randomized initialization. Follow-up methods proposed efficiency improvements via path-space pruning [Li et al. 2022; Wang et al. 2020], guiding [Fan et al. 2023], or reusing neighboring specular paths as seeds [Xu et al. 2023]. But these methods involve expensive data

structures or long initialization times which are problematic for interactive rendering.

Several approaches improve specular path sampling by modifying or replacing the Newton solver. Jhang and Chang [2022] introduce large mutations to correct problematic path configurations to improve Newton solver success rates and Granizo-Hidalgo and Holzschuch [2024a] use the Nelder-Mead algorithm to avoid calculating derivatives. They are orthogonal to our method. Fan et al. [2024] propose iteration-free polynomial solvers, which are almost noisefree but do not scale to longer specular chains. Granizo-Hidalgo and Holzschuch [2024b] reduce the search space dimension but are limited to fast rendering of water caustics.

Interactive caustic rendering. Interactive caustic rendering methods are mostly based on photon mapping and thus inherit its limitations. Earlier work based on caustic mapping [Shah et al. 2007; Wyman 2008] have strong assumptions about receiver geometry. A recent ray-traced method [Ouyang 2023] relies on light-space guiding for noise reduction, limiting supported light types. Noise can also be reduced by combining photon gathering with ReSTIR [Kern et al. 2024], but the bias of photon methods remains. Our method is the first interactive caustic method to employ specular manifold sampling, combining guiding, resampling, and solver modifications in a novel way to enable unbiased caustics at interactive rates.

3 Sample space partitioning

Specular Manifold Sampling [Zeltner et al. 2020] provides an unbiased sampling method for caustic paths, but at a substantial computation cost per individual Monte Carlo sample. Our sample space partitioning method is designed to significantly reduce this cost. In this section, we first describe the necessary details of SMS (Section 3.1) and then present our extensions.

SMS preliminaries 3.1

Let $[x_1, x_2..., x_n]$ represent the vertices along a specular light path. Given a shading point x_1 (usually the first non-specular vertex sampled on the camera subpath) and a sampled emitter vertex \mathbf{x}_n , SMS starts by sampling a seed path $\overline{\mathbf{x}} = [\mathbf{x}_2, ..., \mathbf{x}_{n-1}]$ and then performs a manifold walk [Jakob and Marschner 2012] that iteratively modifies \bar{x} via alternating Newton iterations and ray tracing to find a valid caustic path formed by specular reflections or refractions. We denote a solution path after this process successfully converges as $\overline{\mathbf{x}}^*$. SMS assumes a discrete set of solutions, ignoring symmetric scene constructions with a continuous 1D subspace of solutions.

Because the rest of the vertices are uniquely determined from a sampled \mathbf{x}_2 according to the specular constraints¹, the sample space is only 2D (e.g. representing directions from x_1 to x_2) and can be visualized on a plane (Figure 2). For convenience, the 2D primary sample space $\mathcal{U} = [0, 1)^2$ for \mathbf{x}_2 is used to parameterize the light paths: seed $\overline{\mathbf{x}}$ and the solution $\overline{\mathbf{x}}^*$ can be written as \mathbf{u} and \mathbf{u}^* , respectively, such that $u, u^* \in \mathcal{U}$.

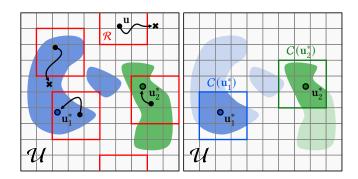


Fig. 2. Illustration of 2D sample space ${\cal U}$ partitioned into $N \times N = 10 \times 10$ tiles. The example has two solutions \mathbf{u}^* (blue and green dots) with their corresponding shaded convergence basins. Left: Example seeds u (black dots) with their 3×3 trust regions $\mathcal{R}(\mathbf{u})$ (red boxes). Manifold walks that step outside are terminated early (crosses). Right: This shrinks the convergence basins to their intersection with the 3×3 conditional sample spaces $C(\mathbf{u}^*)$ and the inverse probability estimate can be restricted to those subregions.

After sampling a solution, the estimate of the caustic irradiance can be expressed as

$$\langle E_{\mathbf{i}}(\mathbf{x}_1) \rangle = \frac{L_{\mathbf{e}}(\mathbf{x}_n) \operatorname{Tr}(\mathbf{x}_1 \leftrightarrow \overline{\mathbf{x}}^* \leftrightarrow \mathbf{x}_n)}{p(\mathbf{x}_n) P(\overline{\mathbf{x}}^*)}, \qquad (1)$$

where $L_{e}(\mathbf{x}_{n})$ is the emitted radiance from \mathbf{x}_{n} , the throughput $Tr(\mathbf{x}_1 \leftrightarrow \overline{\mathbf{x}}^* \leftrightarrow \mathbf{x}_n)$ summarizes reflectance and geometric terms between the vertices, $p(\mathbf{x}_n)$ is the PDF of sampling \mathbf{x}_n , and $P(\overline{\mathbf{x}}^*)$ is the probability of sampling the admissible specular path $\bar{\mathbf{x}}^*$. This path sampling probability equals the area of the convergence basin of the path, the region in 2D primary sample space formed by all seeds **u** that converge to the same caustic path $\overline{\mathbf{x}}^*$:

$$P(\overline{\mathbf{x}}^*) = \int_{\mathcal{U}} p(\mathbf{u}) P(\mathbf{u}^* | \mathbf{u}) d\mathbf{u}, \qquad (2)$$

where $p(\mathbf{u})$ is the PDF for sampling the seed and $P(\mathbf{u}^*|\mathbf{u}) \in \{0, 1\}$ is a binary conditional probability, determining whether the seed **u** converges to $\overline{\mathbf{u}}^*$ via iterations of the manifold walk.

SMS uses a uniform seed PDF $p(\mathbf{u}) = 1$ and estimates the integral $P(\overline{\mathbf{x}}^*)$ numerically using Bernoulli trials [Booth 2007; Qin et al. 2015] to get an unbiased contribution weight (UCW) $W_{\overline{\mathbf{x}}^*}$ such that its expected value is the reciprocal of the probability, i.e. $\mathbb{E}[W_{\overline{\mathbf{x}}^*}] =$ $1/P(\overline{\mathbf{x}}^*)$. More specifically, each Bernoulli trial begins with a random \mathbf{u} and checks if it converges to \mathbf{u}^* using a manifold walk. $W_{\overline{\mathbf{v}}^*}$ is estimated as the number of Bernoulli trials performed until a random \mathbf{u} that converges to \mathbf{u}^* is found. The expected number of required trials is $1/P(\overline{\mathbf{x}}^*)$, which is costly when the probability is small, i.e. for solution paths with small convergence basins. As a result, SMS can involve a large number of Newton iterations, each involving the generation of a new path via ray tracing, making it an expensive sampling technique.

3.2 Motivation

To summarize, SMS consists of two phases: (1) the main sampling phase where the Newton solver is run once to find an admissible path, and (2) the inverse probability estimate using a sequence of Bernoulli trials, each involving another run of the Newton solver.

¹The original implementation of SMS explicitly marks objects as caustic "casters" and sets a fixed n for each object, using deterministic refraction behavior on dieletrics.

Assuming K solutions, the expected number of Newton solver invocations is therefore

$$1 + \sum_{k=1}^{K} P_k \mathbb{E}[\text{\#trials for } k\text{-th solution}] = 1 + \sum_{k=1}^{K} P_k \frac{1}{P_k} = 1 + K, (3)$$

where P_k is the probability of sampling the k-th solution (Eq. 2).

The iterative Newton solver itself quickly becomes the computational bottleneck, especially on GPUs. Each iteration involves computing partial derivatives of the path geometry (for estimating a step towards the solution) and multiple ray tracing operations (to re-project the specular vertices onto the geometry).

Our main optimization goal is to reduce the total number of Newton iterations. We achieve this by introducing *trust regions*, which limit the manifold walks to the vicinity of their starting location. This further allows the inverse probability estimate to be constrained to a *conditional sample space* and requires fewer expected Bernoulli trials.

3.3 Trust regions and conditional sample probabilities

We define trust regions $\mathcal{R}(u) \subset \mathcal{U}$ surrounding all seed paths u. The Newton solver is modified to terminate early in case it steps outside of it, which cuts down the number of Newton iterations for computationally costly cases that would either diverge or only converge slowly².

For a practical implementation, we partition $\mathcal U$ into $N\times N$ regular tiles and define the trust region $\mathcal R(\mathbf u)$ as the 3×3 neighborhood of closest tiles around $\mathbf u$ (with periodic wrapping), see Figure 2 (left). By increasing N we can shrink the trust region size in $\mathcal U$.

Conversely, only seeds that start within that same neighborhood around \mathbf{u}^* can successfully converge and we can limit our search space accordingly. We call this the conditional sample space $C(\mathbf{u}^*)$. In our case, $C(\mathbf{u}^*) = \mathcal{R}(\mathbf{u}^*)$, the 3×3 closest tiles around \mathbf{u}^* . Effectively, the convergence basin of \mathbf{u}^* shrinks based on the intersection with $C(\mathbf{u}^*)$, see Figure 2 (right).

Analogous to Eq. (2), the probability of sampling a given solution with the modified Newton solver is

$$P'(\overline{\mathbf{x}}^*) = \int_{C(\mathbf{u}^*)} p(\mathbf{u}) P_C(\mathbf{u}^* | \mathbf{u}) d\mathbf{u}, \qquad (4)$$

where the binary conditional probability $P_C(\mathbf{u}^*|\mathbf{u})$ now only considers successful manifold walks that stay inside $C(\mathbf{u}^*)$. This can further be factored into

$$P'(\overline{\mathbf{x}}^*) = P(C(\mathbf{u}^*)) \int_{C(\mathbf{u}^*)} \frac{p(\mathbf{u})}{P(C(\mathbf{u}^*))} P_C(\mathbf{u}^*|\mathbf{u}) \, d\mathbf{u} , \qquad (5)$$

with marginal probability $P(C(\mathbf{u}^*)) = \int_{C(\mathbf{u}^*)} p(\mathbf{u}) d\mathbf{u}$. With a uniform seed PDF $p(\mathbf{u})$, this can be evaluated in closed-form with our tiling scheme³, while only the inverse of the conditional probability integral in Eq. (5) remains to be estimated using Bernoulli trials.

This is beneficial because it also scales down the expected number of Bernoulli trials by a factor of $P(C(\mathbf{u}^*))$ and leads to a lower

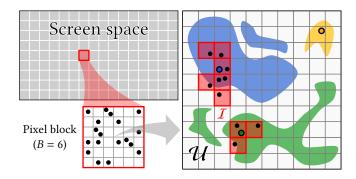


Fig. 3. Illustration of building the prior distributions. **Left**: Screen space is partitioned into $B \times B$ pixel blocks where we select a number of pixels (black dots) to find admissible caustic paths via manifold walks. **Right**: Sample space visualized for the center of the pixel block, with three solutions and their convergence basins marked (blue, green, yellow). Solutions found from within the pixel block (black dots) tend to be spatially coherent and we aggregate the union of all corresponding sample space tiles as $I \subset \mathcal{U}$. During rendering we will place samples inside I with higher probability.

number of expected Newton solver invocations, compared to Eq. (3):

$$1 + \sum_{k=1}^{K} P_k' \mathbb{E}[\text{\#trials for } k\text{-th sol.}] = 1 + \sum_{k=1}^{K} P_k' \frac{P(C_k)}{P_k'} < 1 + K . \quad (6)$$

 P_k' is the probability of sampling the k-th solution (Eq. (4)) and $P(C_k)$ is a shorthand for the marginal probability of that solution.

3.4 Trade-off between variance and performance

All modifications so far were designed to improve the performance of SMS, by reducing both the number of Newton iterations and Bernoulli trials. However, decreasing the size of trust regions (i.e., choosing large N) lowers the Newton solver success rate for finding admissible paths compared to the original SMS (i.e., using N = 1).

To make SMS effective for interactive rendering, we need to strike a balance between the individual sample cost and the expected variance. The next section discusses how to improve both aspects at the same time, using prior knowledge about the sample locations.

4 Per-frame prior distributions

We will now counteract the variance increase from Section 3 by constructing a simple importance sampling distribution on top of the sample space partitioning scheme. We start with the simplified case of a single specular object and a point light source (Section 4.1) before extending the concept to more general cases (Section 4.2).

4.1 Building and sampling from prior distributions

Before rendering each frame we run a brief precomputation step to identify regions of sample space where solutions are most likely to occur. As observed by Xu et al. [2023], solutions and their convergence basins are highly correlated for nearby shading points; we exploit this by sharing information between pixels in screen space.

For each $B \times B$ pixel block (B = 16 in our implementation), we map threads to random pixels and use manifold walks—unrestricted by trust regions in this case—to try and find admissible caustic paths

²Ideally, the mapping between sample and world space is continuous. This is true when sampling vertices based on directions but not necessarily with surface-based sampling. ³Other trust region constructions and seed PDFs are possible as well, as long as we can evaluate the integral and draw samples proportional to the density.

from the corresponding shading points to the light source. We mark all sample space tiles where solutions were found as important and store their union $\mathcal{I} \subset \mathcal{U}$ per block, see Figure 3. We use 128 threads by default and increase that to 256 for more difficult cases (Figure 1).

When rendering the actual frame, we sample seed paths inside I more often, which increases the likelihood of converging to a solution. Our new seed path PDF is

$$p(\mathbf{u}) = \alpha p_T(\mathbf{u}) + (1 - \alpha) p_U(\mathbf{u}), \qquad (7)$$

where the uniform density $p_{\mathcal{I}}(\mathbf{u})$ in \mathcal{I} is mixed with the uniform density $p_{\mathcal{U}}(\mathbf{u})$ in \mathcal{U} to ensure unbiased coverage of all possible solutions. We use $\alpha = 0.8$ as a mixing weight in practice.

We experimented with additionally weighting the PDF of tiles in I by the irradiance of their corresponding solution. However, irradiance estimates require Bernoulli trials and the added cost during the precomputation step did not translate into higher sampling efficiency overall.

Importance sampling is applied on top of the conditional sample space concept from Section 3.3. Eq. (5) still holds with the new seed path PDF (Eq. (7)) and the marginal probability integral still has a simple closed-form solution

$$P(C(\mathbf{u}^*)) = \alpha \frac{\#(I \cap C(\mathbf{u}^*))}{\#(I)} + (1 - \alpha) \frac{M^2}{N^2},$$
 (8)

where $\#(\cdot)$ is the number of tiles in the region.

Samples drawn for the Bernoulli trials when estimating the inverse probability are also still restricted based on the conditional sample space: they are sampled uniformly, either inside the intersection of I and $C(\mathbf{u}^*)$, or the entirety of $C(\mathbf{u}^*)$ based on an adjusted mixing weight

$$\alpha' = \alpha \frac{\#(I \cap C(\mathbf{u}^*))}{\#(I)} \frac{1}{P(C(\mathbf{u}^*))}. \tag{9}$$

Extension to multiple lights and specular objects

While the precomputation from Section 4.1 could be repeated for all pairs of light sources and specular objects in a given scene, the computational cost and memory requirement would scale poorly. Instead, we keep the precomputation budget fixed but allow each sample inside the pixel block to pick a random light-object pair, as well as a position \mathbf{x}_n on that light source. We can then aggregate important tiles that were found for the same light-object pairs, which we store in linked lists indexed by the unique light-object pair⁴.

At any given shading point during rendering, we then either pick from this set of pairs (with probability α), or randomly from the whole scene (otherwise), followed by picking the position on the light. Conditioned on all of these, we then finally try sampling an admissible caustic path connecting the two.

4.3 Discussion

The combination of our tiled sample space partitioning and prior distribution improves both performance and image quality compared to the original SMS. Importance sampled seeds tend to start close to solutions, requiring fewer Newton iterations and Bernoulli

trials. The smaller convergence basins due to the trust regions also avoid a lot of wasted computation in scene regions without any caustics, similar to the selective activation in Fan et al. [2023].

Overall we mitigate the primary source of fireflies in caustic sampling while keeping a low construction overhead. Residual noise remains due to (1) limited resolution and accuracy of the prior distribution, and (2) not considering the emitted energy of different light sources. As a last step, we incorporate spatiotemporal resampling to get closer to the ideal target distribution in practice.

Spatiotemporal Resampling

We have combined SMS with sample space partitioning to obtain an efficient caustic sampler and, while it does not perform perfect importance sampling, it effectively avoids high-energy fireflies. This makes the generated samples good candidates to feed into a resampling algorithm that further improves the distribution. The similarity of specular paths across nearby pixels [Xu et al. 2023] also allows us to leverage the power of ReSTIR [Bitterli et al. 2020] to harvest a large number of candidate samples with spatiotemporal reuse.

GRIS preliminaries

Generalized Resampled Importance Sampling (GRIS) [Lin et al. 2022] takes candidate samples X_i from different domains Ω_i and maps them to the target domain Ω with bijective shift mapping functions $T_i: \Omega_i \to \Omega$. The goal is to resample from the candidates to move the sample distribution towards a target distribution proportional to a function $\hat{p}(y)$. By computing the shift mapping $Y_i = T_i(X_i)$, the target function value $\hat{p}(Y_i)$, the Jacobian determinant of the shift mapping $|\partial T_i/\partial X_i|$, and incorporating the UCW of X_i , the resampling weights can be evaluated as

$$w_i = m_i(Y_i)\hat{p}(Y_i)W_{X_i} \left| \frac{\partial T_i}{\partial X_i} \right| . \tag{10}$$

Here, $m_i(y)$ are resampling MIS weights ensuring coverage over Ω with $\sum_{i=1}^{M} [y \in T_i(\text{supp } X_i)] m_i(y) = 1$. A sample Y is then chosen from Y_i with probabilities proportional to w_i , with UCW

$$W_Y = \frac{1}{\hat{p}(Y)} \sum_{i=1}^{M} w_i . {11}$$

As $Var[\sum_i w_i] \to 0$, the distribution of Y moves towards \hat{p} , and $f(Y)W_Y$ yields a low-variance estimate of $\int_{\Omega} f(y) dy$ (usually a path integral) for $f \approx \hat{p}$.

GRIS can be iteratively chained, reusing samples across pixels and frames—an idea underpinning ReSTIR, where each pixel maintains a reservoir (X, W_X, c) that carries information required for resampling. A rendered frame begins with path-traced samples, followed by phases of temporal and then spatial resampling, each updating the reservoirs. These are then used both for shading and as candidate samples for the next frame via motion vectors. The confidence weight c (initialized as 1) is updated by summing the c values of input reservoirs used as candidates. It scales the importance of a domain in m_i similar to multi-sample MIS, but is capped to limit unconstrained temporal correlation.

 $^{^4}$ The previous assumption of coherent solutions in ${\cal U}$ only holds for small light sources. To support large area lights or environment maps, they would need to be subdivided into multiple lights with smaller coherent extents.

While ReSTIR improves sampling efficiency, its convergence is hindered if the initial sampling distribution is poor, or the shift mappings are inaccurate. In this work, we enhance both: we generate high-quality initial samples and apply manifold shift mappings [Jakob and Marschner 2012; Lehtinen et al. 2013] in GRIS, leveraging spatiotemporal coherence of solutions. The resulting technique is more effective at rendering caustics than prior methods, e.g., ReSTIR PT [Lin et al. 2022], that use path tracing for the initial samples and rely on random replay to shift caustic paths.

5.2 GRIS for SMS paths

Focusing on caustic sampling, we define the domain Ω of interest of a pixel to be all specular paths that can be sampled by SMS. The samples $\hat{\mathbf{x}} = [\overline{\mathbf{x}}^*, \mathbf{x}_n]$ contain a vertex $\mathbf{x}_n \in \mathcal{L}$ on the set of all emissive surfaces, and an admissible specular subpath $\overline{\mathbf{x}}^* \in \mathcal{S}(\mathbf{x}_1, \mathbf{x}_n)$ from the discrete set of solutions connecting the two path endpoints. For simplicity, we consider the case of perfectly specular delta BSDFs, i.e. with zero roughness. With the notation in Section 3.1 and defining ρ as the BSDF at \mathbf{x}_1 , the integrand we want to solve is

$$L(\mathbf{x}_1, \omega_0) = \int_{\mathcal{L}} \left(\sum_{\overline{\mathbf{x}}^* \in \mathcal{S}(\mathbf{x}_1, \mathbf{x}_n)} f(\hat{\mathbf{x}}) \right) d\mathbf{x}_n$$
 (12)

with

$$f(\hat{\mathbf{x}}) = \rho(\mathbf{x}_1, \omega_0, \omega_i^*) \operatorname{Tr}(\mathbf{x}_1 \leftrightarrow \overline{\mathbf{x}}^* \leftrightarrow \mathbf{x}_n) L_{\mathbf{e}}(\mathbf{x}_n) . \tag{13}$$

Eq. (12) integrates over all emissive points \mathbf{x}_n and accumulates the energy from all discrete caustic solutions towards \mathbf{x}_1 .

To estimate this integral, we sample an admissible specular path $\hat{\mathbf{x}} \in \Omega$ with our technique from Sections 3 and 4 before computing its UCW $W_{\hat{\mathbf{x}}} = W_{\overline{\mathbf{x}}^*}/p(\mathbf{x}_n)$ with Bernoulli trials.

We follow the common approach [Lin et al. 2022; Ouyang et al. 2021] to define $\hat{p}(\hat{\mathbf{x}}) = f(\hat{\mathbf{x}})$. We apply manifold shift mappings [Lehtinen et al. 2013] to reuse samples from other pixels (domains). This shift mapping $T_i(\overline{\mathbf{x}}^*, \mathbf{x}_n) = (\overline{\mathbf{x}}^*, \mathbf{x}_n)$ preserves the emissive vertex \mathbf{x}_n from the base path in Ω_i^5 and shifts the specular subpath according to the changed shading point \mathbf{x}_1 . This is accomplished with a manifold walk, starting at seed $\overline{\mathbf{x}}^*$, that tries to find a solution in the current pixel. In case it successfully converges to a solution $\overline{\mathbf{x}}^{*\prime}$ in Ω , we also need to perform the inverse shift mapping T_i^{-1} to verify invertibility to the original $\overline{\mathbf{x}}^*$. Since the shift mapping does not involve a density change⁶, the Jacobian determinant of the mapping is 1.

For the resampling MIS weights m_i , we choose the generalized pairwise heuristic as recommended by Lin et al. [2022], which requires additionally shifting the sample from the current pixel to each neighboring pixel i. This results in a total of four manifold walks for each neighbor pixel we reuse from.

For the general case of glossy non-delta BSDFs, we follow prior work by extending our sampling technique to *offset manifolds* [Jakob and Marschner 2012]. This replaces the shading normal in the specular constraint with a sampled half vector and updates the corresponding tangent space. To reuse a path, we shift its offset manifold

first by replaying the random numbers associated with the half vector generation, before performing a manifold walk on the shifted offset manifold. By sampling the normal distribution function (NDF), this generates the same half vectors as the base path, but requires only storing the RNG seed. The Jacobian determinant of this shift in area measure is not 1, but can be easily computed [Lin et al. 2022].

Our method can also handle indirect caustics, e.g. in SDS paths. Non-delta vertices in the prefix path before \mathbf{x}_1 can either be shifted using random replay as in ReSTIR PT [Lin et al. 2022], or kept independent as in Suffix ReSTIR [Kettunen et al. 2023]. The manifold shift is then applied only in the conditional space created by \mathbf{x}_1 .

5.3 Performance of manifold shift mappings

Prior work in gradient domain path tracing [Hua et al. 2019; Kettunen et al. 2015] avoided manifold shift mappings due to their cost and instead opted for simpler ones like half-vector copy or random replay. This assessment, however, needs to be re-evaluated in the context of SMS and ReSTIR. First, shifts like half-vector copy are prone to fail for difficult scenes (e.g., scenes with tiny light sources) where SMS shows its advantage. Second, each reuse in ReSTIR involves only a constant number of manifold walks—but effectively inherits many historical samples that do not require running more manifold walks for their generation or inverse probability estimate. This latter behavior is actually similar to the original manifold exploration technique [Jakob and Marschner 2012], a result of the similarity between ReSTIR and MCMC [Sawhney et al. 2024].

To further optimize the performance of resampling, we borrow ideas from our trust region method. Although, instead of enforcing a trust region, we simply limit the maximum number of Newton iterations for shift mappings to a small number (we use 5 in our results). This avoids wasted computation for mappings between distant specular paths that are prone to failure, and remains unbiased.

6 Results

We implemented our method in the Falcor framework [Kallweit et al. 2022] and ran the experiments on a PC with an i9-13900K CPU and RTX 4090 GPU. We test with the Reflective Plane, the Double-Refractive Slab, and the Swimming Pool scenes from Zeltner et al. [2020]. The Reflective Plane and the Double-Refractive Slab are lit with a small area light while the Swimming Pool uses a high-frequency environment light. We first compare our method against prior work, then we validate its building blocks. All images are rendered at 1920×1080 resolution and we report RMSE compared to high-SPP path traced references.

Comparisons against previous methods. Figure 1 shows results on a composite scene that contains both the plane and the slab and includes two area lights. Our method outperforms SMS in both speed and quality by leveraging trust regions and prior distributions. Combined with ReSTIR, our method delivers high-quality caustics at only 1 SPP by aggregating samples from different light sources, shapes, and specular solutions.

Figure 4 shows rendering comparisons of our method against PT and ReSTIR PT [Lin et al. 2022]. Our method uses 1 SPP to produce significantly better results, and ReSTIR further improves our results while adding a small amount of overhead. With higher

 $^{^5\}mathrm{We}$ reuse the light space coordinates to allow temporal reuse with moving lights.

⁶We assume light sources do not change their areas over time.

Table 1. Statistics and timings for different stages of our algorithm on scenes in Figures 1 and 4 ("Ours + ReSTIR" has the additional resampling stage). In the manifold sampling stage, we provide a comparison between our sample space partitioning and the original SMS.

	Prior Construction (Ours)			Manifold Sampling (SMS / Ours)					Resampling ("Ours+ReSTIR")	
	Sample Success Rate	Newtons Per Sample	Time (ms)	Sample Success Rate	Newtons Per Sample	Bernoulli Per Solution	Newtons Per Solution	Time (ms)	Temporal Time (ms)	Spatial Time (ms)
Double-Refractive Slab	11.7%	11.08	3.43	9.4% / 21.6%	8.76 / 1.82	4.24 / 1.43	67.42 / 8.67	38.41 / 12.42	2.53	4.60
REFLECTIVE PLANE	8.0%	19.16	1.71	8.1% / 19.4%	19.15 / 4.85	66.9 / 9.25	1182.72 / 39.05	120.30 / 26.00	2.82	3.92
SWIMMING POOL	12.4%	5.67	0.72	10.8% / 34.7%	5.73 / 5.30	6.10 / 2.02	42.67 / 17.03	8.82 / 9.96	1.68	2.07
PLANE AND SLAB (Figure 1)	11.7%	16.17	10.62	10.5% / 34.7%	15.78 / 2.67	21.09 / 1.75	144.92 / 9.05	117.02 / 41.69	4.42	20.02

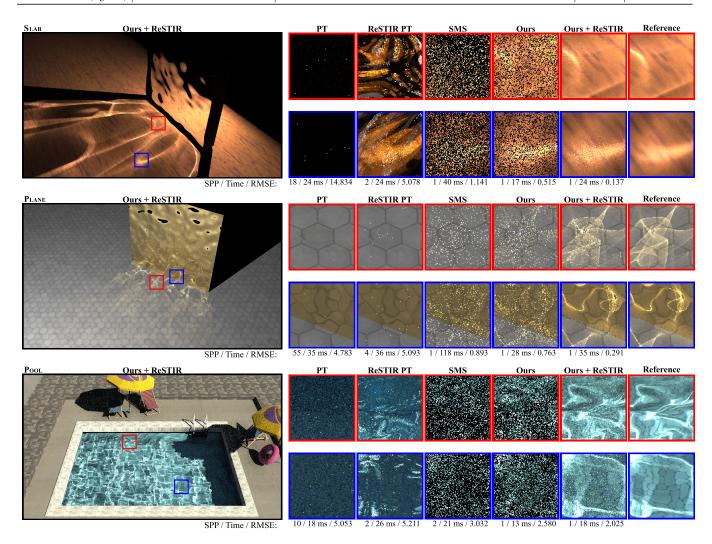


Fig. 4. Comparison between our method and prior work in two different scenes. The reference is generated with path tracing at high SPP. For "Ours + ReSTIR", we use the ReSTIR parameters recommended by Wyman and Panteleev [2021], setting the maximum cap for temporal reuse at c = 20, and randomly reusing one spatial neighbor in a 30-pixel radius to dither out temporal correlation.

frame times, SMS exhibits higher variance; Path Tracing struggles finding any solutions via naive BSDF sampling; and ReSTIR PT spreads the small number of outlier paths found by path tracing spatiotemporally to create distracting correlation artifacts. For our

method, we partition the sample space with 16×16 tiles (N = 16) and set the trust region size to be 3×3 tiles. For both SMS and our method, we set the maximum number of Bernoulli trials as 128. This

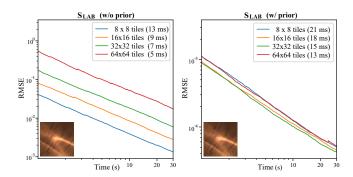


Fig. 5. Convergence plots for our method without (left) and with (right) prior distributions. Each case is tested with 4 varying tiling resolutions N. The RMSE numbers are computed on a rendered crop of the caustic from the Slab scene, accumulated at different time stamps.

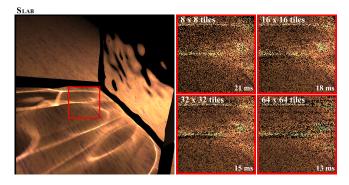


Fig. 6. When keeping a fixed trust region size of 3×3 , increasing the tiling resolution makes the trust region smaller, which brings better quality for coherent regions (solutions are close to each other for nearby shading points) while reducing the quality in incoherent regions.

clamping is theoretically biased, but it is enough for our method to deliver unbiased quality in our tests.

Prior distributions parameters. Figure 5 shows ablations with and without our prior distributions at different tiling resolutions N. Larger N partitions the sample space into more tiles and shrinks the trust regions. Without the prior distribution (left plot), larger N accelerates performance at 1 SPP but increases variance and slows convergence. With the prior distribution (right plot), the frame times increase, but the variance reduction is significant. The optimal convergence speed occurs at N=16 and N=32, which we use in the other results.

Figure 6 shows the visual effects of the tiling resolution N on the prior distributions. A smaller trust region (larger N) generally improves performance and can help to specify the location of solutions more accurately, potentially improving the quality. In contrast, increasing N without the prior distribution only increases variance. However, if the trust regions become too small, important tiles found by neighboring pixels cannot be reused effectively, decreasing the sampling efficiency. This effect is especially severe in cases where solutions among nearby pixels vary sharply.

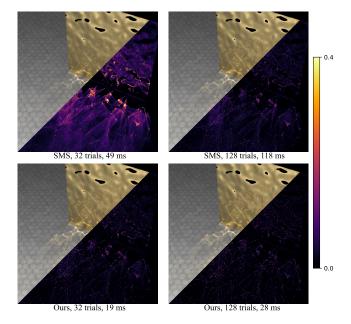


Fig. 7. False-color visualization of the pixel-wise absolute error using different maximum Bernoulli trials. 5000 frames are accumulated to judge the converged quality. The corresponding references are computed using SMS with 1024 trials and accumulating 10000 frames. This shows our tiling technique can achieve unbiased results with a smaller maximum number of Bernoulli trials while insufficient Bernoulli trials bring energy loss.

Trust regions reduce Bernoulli trials. Figure 7 visualizes the energy loss given a fixed Bernoulli trial budget for our method and SMS. An insufficient number of Bernoulli trials brings energy loss. Our trust region generally yields a larger conditional probability and requires less Bernoulli trials to estimate the inverse probability. Therefore, our method reaches unbiased results more quickly compared to SMS in practice.

Validation of ReSTIR. We utilize ReSTIR to aggregate spatiotemporal samples. ReSTIR provides a dramatic overall noise reduction, which benefits from the temporal reuse. However, for incoherent regions where the solution paths of neighboring pixels differ a lot, correlation artifacts (in the form of blocky noise patterns) can arise due to inaccurate motion vectors. In these cases, spatial reuse is useful to reduce correlation and variance (Figure 8).

Comparisons to ReSTIR-FG. ReSTIR-FG [Kern et al. 2024] reuses temporal caustic photons within a given radius for density estimation without applying shift mappings. As a result, reuse fails when no temporal photons fall within the neighborhood of the current shading point. Moreover, the lack of bijective shift mappings leads to visible lagging artifacts in dynamic scenes. Our method achieves more robust temporal reuse by doing bijective shift mappings. Figure 9 presents an equal-time comparison between our method and ReSTIR-FG under camera motion.

Statistic Analysis. Table 1 reports statistics and timings for the test scenes shown in Figures 1 and 4. Compared to SMS, our method increases the rate of successfully sampling a solution by introduc-

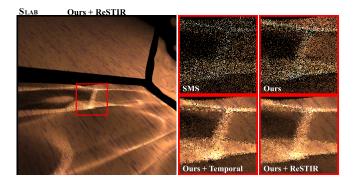


Fig. 8. A comparison between SMS, our method, and two ReSTIR variants used in conjunction with our method. With only temporal resampling ("Ours + Temporal") noise reduction is significant overall. But difficult regions (e.g., "edges" of caustics) cause visible correlation artifacts. Adding spatial reuse ("Ours + ReSTIR") lowers the variance further and improves the noise characteristics.

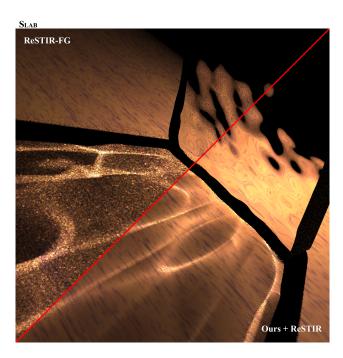


Fig. 9. An equal-time (24 ms) comparison between our method and ReSTIR-FG under camera motion. ReSTIR-FG traces 10 million photons per frame but fails to reuse temporal caustic photons beyond its radius, leading to visible temporal artifacts. In contrast, our method employs shift mappings with bijectivity checks, enabling more robust temporal reuse.

ing prior distributions, which in turn improves image quality. It also reduces the average number of Newton iterations per SMS sample (including failed samples) by constraining the solver, and lowers per-solution cost by decreasing Bernoulli trials via inverse probability estimation in the conditional space. In the SWIMMING POOL, our method is slightly slower than SMS because it samples more solutions and needs more overall Newton iterations. We also

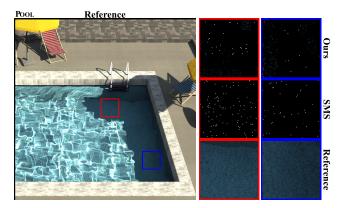


Fig. 10. A failure case where our prior distribution fails to provide useful information for the sampled light sources. In this situation, the sampling quality may be worse than SMS, due to the reduced probability of sampling valid solutions

noticed that thread divergence becomes larger when there are multiple objects with different numbers of bounces in Figure 1, which significantly increases the cost of each stage.

Limitations. Our method reverses the sampling quality reduction from limited Newton iterations by introducing priors. However, if the prior distribution fails to cover a sampled light-object pair due to insufficient samples, the probability of finding a solution drops and the variance increases. Figure 10 illustrates such a case, where the sampled environment light texels are not covered by the prior distribution. The shadowed region is mainly lit by low-frequency texels of the environment map, but these texels are undersampled during prior construction. In such cases, we fall back to sample-space partitioning without priors, which has a lower solution probability and thus higher variance.

Conclusion and future work

We introduced a sample space partitioning method that simultaneously improves runtime cost and quality of Specular Manifold Sampling. By restricting the Newton solver to a vicinity of its starting location, and importance sampling these regions using a prior distribution, our method effectively culls unnecessary computation. We share the prior distribution among nearby image pixels to amortize the cost of generating the distribution. And we take the reuse concept further by incorporating spatiotemporal resampling (ReSTIR), aggregating samples over space and time to reduce the variance to a very low level, producing high-quality caustics in real time. As an improvement of unbiased SMS, our method does not suffer from energy loss and can use 1 SPP inputs for low-cost sample generation, unlike prior work [Xu et al. 2023] that presamples initial guesses to seed biased SMS.

We favor speed over quality in building prior distribution, ignoring solution energy to avoid Bernoulli trials in presampling. Future work may explore ways to build better prior distributions at a similar low cost. Choosing the size of the trust regions adaptively, or at non-uniform resolution across the scene, could also lead to further benefits.

Performance-wise, the repeated ray tracing during the manifold walks still takes a lot of time and thread divergence can still be large due to different numbers of Newton iterations and Bernoulli trials required for the random sampling happening for pixels in the same warp. Future work may explore better surface parameterizations so that manifold walks can directly operate in a space that does not require re-projections via ray tracing. A more consistent UV parameterization would also improve the efficiency of our prior distribution. Lastly, how to better schedule the workload of SMS with ReSTIR on GPUs to relieve thread divergence and boost performance is an interesting future direction as well.

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 Trans. Graph. 39, 4, Article 148 (Aug. 2020), 17 pages. doi:10.1145/3386569.3392481
- Thomas E. Booth. 2007. Unbiased Monte Carlo Estimation of the Reciprocal of an Integral. Nuclear Science and Engineering 156 (2007), 403 407.
- Brent Burley, David Adler, Matt Jen-Yuan Chiang, Hank Driskill, Ralf Habel, Patrick Kelly, Peter Kutz, Yining Karl Li, and Daniel Teece. 2018. The design and evolution of disney's hyperion renderer. ACM Transactions on Graphics (TOG) 37, 3 (2018), 1–22.
- D5Render. 2023. How to Create Caustics in Real-time Rendering. https://www.d5render.com/posts/caustics-in-render Blog post, D5 Render.
- Honghao Dong, Guoping Wang, and Sheng Li. 2023. Neural parametric mixtures for path guiding. In ACM SIGGRAPH 2023 Conference Proceedings. 1–10.
- Zhimin Fan, Jie Guo, Yiming Wang, Tianyu Xiao, Hao Zhang, Chenxi Zhou, Zhenyu Chen, Pengpei Hong, Yanwen Guo, and Ling-Qi Yan. 2024. Specular Polynomials. *ACM Transactions on Graphics (TOG)* 43, 4 (2024), 1–13.
- Zhimin Fan, Pengpei Hong, Jie Guo, Changqing Zou, Yanwen Guo, and Ling-Qi Yan. 2023. Manifold path guiding for importance sampling specular chains. ACM Transactions on Graphics (TOG) 42, 6 (2023), 1–14.
- Luca Fascione, Johannes Hanika, Mark Leone, Marc Droske, Jorge Schwarzhaupt, Tomáš Davidovič, Andrea Weidlich, and Johannes Meng. 2018. Manuka: A batch-shading architecture for spectral path tracing in movie production. ACM Transactions on Graphics (TOG) 37, 3 (2018), 1–18.
- Iliyan Georgiev, Jaroslav Krivánek, Tomas Davidovic, and Philipp Slusallek. 2012. Light transport simulation with vertex connection and merging. ACM Trans. Graph. 31, 6 (2012), 192–1.
- Ana Granizo-Hidalgo and Nicolas Holzschuch. 2024a. Computing Manifold Next-Event Estimation without Derivatives using the Nelder-Mead Method. In Eurographics Symposium on Rendering, Eric Haines and Elena Garces (Eds.). The Eurographics Association. doi:10.2312/sr.20241156
- Ana Granizo-Hidalgo and Nicolas Holzschuch. 2024b. Interactive Rendering of Caustics using Dimension Reduction for Manifold Next-Event Estimation. *Proceedings of the ACM on Computer Graphics and Interactive Techniques* 7, 1 (2024), 1–16.
- Toshiya Hachisuka and Henrik Wann Jensen. 2009. Stochastic Progressive Photon Mapping. ACM Transactions on Graphics 28, 5 (2009).
- Toshiya Hachisuka, Anton S Kaplanyan, and Carsten Dachsbacher. 2014. Multiplexed metropolis light transport. ACM Transactions on Graphics (TOG) 33, 4 (2014), 1–10.
- Johannes Hanika, Marc Droske, and Luca Fascione. 2015. Manifold Next Event Estimation. Computer Graphics Forum (2015).
- Binh-Son Hua, Adrien Gruson, Victor Petitjean, Matthias Zwicker, Derek Nowrouzezahrai, Elmar Eisemann, and Toshiya Hachisuka. 2019. A Survey on Gradient-Domain Rendering. In Computer Graphics Forum, Vol. 38. Wiley Online Library. 455–472.
- Wenzel Jakob and Steve Marschner. 2012. Manifold Exploration: A Markov Chain Monte Carlo Technique for Rendering Scenes with Difficult Specular Transport. ACM Trans. Graph. 31, 4 (July 2012), 58:1–58:13.
- Henrik Wann Jensen. 1996. Global illumination using photon maps. In Rendering Techniques' 96: Proceedings of the Eurographics Workshop in Porto, Portugal, June 17–19, 1996 7. Springer, 21–30.
- Jia-Wun Jhang and Chun-Fa Chang. 2022. Specular manifold bisection sampling for caustics rendering. In Computer Graphics Forum, Vol. 41. Wiley Online Library, 247–254.
- Simon Kallweit, Petrik Clarberg, Craig Kolb, Tom'aš Davidovič, 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.
- Anton S. Kaplanyan, Johannes Hanika, and Carsten Dachsbacher. 2014. The Natural-Constraint Representation of the Path Space for Efficient Light Transport Simulation.

- ACM Transactions on Graphics 33, 4 (July 2014).
- Csaba Kelemen, László Szirmay-Kalos, György Antal, and Ferenc Csonka. 2002. A simple and robust mutation strategy for the metropolis light transport algorithm. In Computer Graphics Forum, Vol. 21. Wiley Online Library, 531–540.
- René Kern, Felix Brüll, and Thorsten Grosch. 2024. ReSTIR FG: Real-Time Reservoir Resampled Photon Final Gathering. In Eurographics Symposium on Rendering, Eric Haines and Elena Garces (Eds.). The Eurographics Association. doi:10.2312/sr. 20241155
- Markus Kettunen, Daqi Lin, Ravi Ramamoorthi, Thomas Bashford-Rogers, and Chris Wyman. 2023. Conditional Resampled Importance Sampling and ReSTIR. 1–11. doi:10.1145/3610548.3618245
- Markus Kettunen, Marco Manzi, Miika Aittala, Jaakko Lehtinen, Frédo Durand, and Matthias Zwicker. 2015. Gradient-domain path tracing. ACM Transactions on Graphics (TOG) 34, 4 (2015), 1–13.
- Jaakko Lehtinen, Tero Karras, Samuli Laine, Miika Aittala, Frédo Durand, and Timo Aila. 2013. Gradient-domain metropolis light transport. ACM Transactions on Graphics (TOG) 32, 4 (2013), 1–12.
- He Li, Beibei Wang, Changhe Tu, Kun Xu, Nicolas Holzschuch, and Ling-Qi Yan. 2022. Unbiased caustics rendering guided by representative specular paths. In SIGGRAPH Asia 2022 Conference Papers. 1–8.
- Daqi Lin, Markus Kettunen, Benedikt Bitterli, Jacopo Pantaleoni, Cem Yuksel, and Chris Wyman. 2022. Generalized resampled importance sampling: foundations of ReSTIR. ACM Trans. Graph. 41, 4, Article 75 (July 2022), 23 pages. doi:10.1145/ 3528223.3530158
- Thomas Müller, Markus Gross, and Jan Novák. 2017. Practical Path Guiding for Efficient Light-Transport Simulation. *Computer Graphics Forum (Proceedings of EGSR)* 36, 4 (June 2017), 91–100.
- Thomas Müller, Brian McWilliams, Fabrice Rousselle, Markus Gross, and Jan Nov'ak. 2018. Neural Importance Sampling. arXiv preprint arXiv:1808.03856 (2018).
- Yaobin Ouyang. 2023. Generating ray-traced caustic effects in Unreal Engine 4, part

 1. https://developer.nvidia.com/blog/generating-ray-traced-caustic-effects-in-unreal-engine-4-part-1/
- Yaobin Ouyang, Shiqiu Liu, Markus Kettunen, Matt Pharr, and Jacopo Pantaleoni. 2021. ReSTIR GI: Path Resampling for Real-Time Path Tracing. Computer Graphics Forum (2021). doi:10.1111/cgf.14378
- Hao Qin, Xin Sun, Qiming Hou, Baining Guo, and Kun Zhou. 2015. Unbiased photon gathering for light transport simulation. ACM Trans. Graph. 34, 6, Article 208 (Nov. 2015), 14 pages. doi:10.1145/2816795.2818119
- Lukas Ruppert, Sebastian Herholz, and Hendrik PA Lensch. 2020. Robust fitting of parallax-aware mixtures for path guiding. ACM Trans. Graph. 39, 4 (2020), 147–1.
- Rohan Sawhney, Daqi Lin, Markus Kettunen, Benedikt Bitterli, Ravi Ramamoorthi, Chris Wyman, and Matt Pharr. 2024. Decorrelating restir samplers via mcmc mutations. ACM Transactions on Graphics 43, 1 (2024), 1–15.
- Musawir A Shah, Jaakko Konttinen, and Sumanta Pattanaik. 2007. Caustics mapping: An image-space technique for real-time caustics. IEEE Transactions on Visualization and Computer Graphics 13, 2 (2007), 272–280.
- Eric Veach and Leonidas Guibas. 1994. Bidirectional Estimators for Light Transport. In Fifth Eurographics Workshop on Rendering. Darmstadt, Germany, 147–162.
- Eric Veach and Leonidas J. Guibas. 1997. Metropolis Light Transport. In Computer Graphics (ACM SIGGRAPH '97 Proceedings), Vol. 31. 65–76.
- Jiří Vorba, Ondřej Karlík, Martin Šik, Tobias Ritschel, and Jaroslav Křivánek. 2014.
 On-line Learning of Parametric Mixture Models for Light Transport Simulation.
 ACM Transactions on Graphics (Proceedings of SIGGRAPH 2014) 33, 4 (aug 2014).
- Beibei Wang, Miloš Hašan, and Ling-Qi Yan. 2020. Path Cuts: Efficient Rendering of Pure Specular Light Transport. ACM Trans. Graph. 39, 6, Article 238 (Nov. 2020), 12 pages.
- Chris Wyman. 2008. Hierarchical caustic maps. In Proceedings of the 2008 symposium on Interactive 3D graphics and games. 163–171.
- Chris Wyman and Alexey Panteleev. 2021. Rearchitecting Spatiotemporal Resampling for Production. In ACM/EG Symposium on High Perfrormance Graphics. 23–41. doi:10. 2312/hpg.20211281
- Xiaofeng Xu, Lu Wang, and Beibei Wang. 2023. Efficient Caustics Rendering via Spatial and Temporal Path Reuse. Computer Graphics Forum (2023). doi:10.1111/cgf.14975
- Tizian Zeltner, Iliyan Georgiev, and Wenzel Jakob. 2020. Specular Manifold Sampling for Rendering High-Frequency Caustics and Glints. ACM Trans. Graph. 39, 4, Article 149 (jul 2020), 15 pages.