




Efficient Caustics Rendering via Spatial and Temporal Path Reuse

Xiaofeng Xu¹  Lu Wang^{†1}  Beibei Wang^{†2,3} 

¹Shandong University, China

²Nankai University, China

³Nanjing University of Science and Technology, China

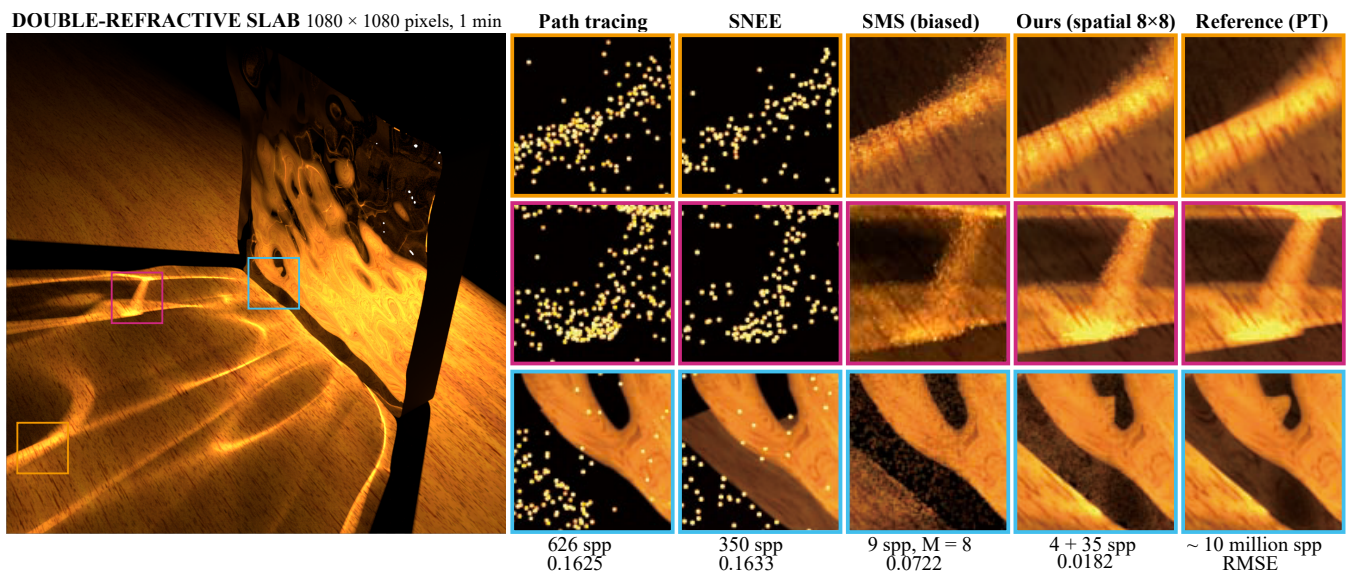


Figure 1: Equal-time comparison (1 minute) among our method, path tracing (PT), specular next event estimation (SNEE) [LZHJ20], biased specular manifold sampling (SMS) [ZGJ20] on the Slab scene. We report the number of samples per pixel (SPP) computed by each method and root mean square error (RMSE) compared to the reference image. Our method employs 4 SPP in the path sample generation phase, with M set to the same value as in the biased SMS method. Our method exhibits the lowest variance both visually and quantitatively.

Abstract

Caustics are complex optical effects caused by the light being concentrated in a small area due to reflection or refraction on surfaces with low roughness, typically under a sharp light source. Rendering caustic effects is challenging for Monte Carlo-based approaches, due to the difficulties of sampling the specular paths. One effective solution is using the specular manifold to locate these valid specular paths. Unfortunately, it needs many iterations to find these paths, leading to a long rendering time. To address this issue, our key insight is that the specular paths tend to be similar for neighboring shading points. To this end, we propose to reuse the specular paths spatially. More specifically, we generate some specular path samples with a low sample rate and then reuse these specular path samples as the initialization for specular manifold walk among neighboring shading points. In this way, much fewer specular path-searching iterations are performed, due to the efficient initialization close to the final solution. Furthermore, this reuse strategy can be extended for dynamic scenes in a temporal manner, such as light moving or specular geometry deformation. Our method outperforms current state-of-the-art methods and can handle multiple bounces of light and various scenes.

CCS Concepts

• Computing methodologies → Rendering;

[†] Co-corresponding authors: L. Wang (luwang_hcivr@sdu.edu.cn) and B. Wang (beibei.wang@njust.edu.cn)

1. Introduction

Scattering from specular surfaces produces complex optical effects, such as intricate caustics due to focused reflection and multiple

refraction. Rendering caustics is challenging because they correspond to high luminous intensity focused over a small area. Current light transport methods are difficult to find the specular paths that contribute to this small area, resulting in poor convergence or large variance.

Path tracing [Kaj86] is an elegant solution to the rendering equation because it considers all types of light transport within the same framework. However, it introduces excessive variance when rendering caustics due to inefficient path sampling between reflective or refractive specular events. At specular interfaces, light must adhere to the laws of reflection or refraction (Snell's law), significantly reducing the probability of sampling an admissible specular path that connects the camera to a light source. To solve this problem, Zeltner et al. [ZGJ20] propose Specular Manifold Sampling (SMS) to sample specular paths. Their method combines deterministic root finding with stochastic sampling in a Monte Carlo framework. The SMS method deals with multiple specular paths and requires a large number of Newton iterations for the root-finding process. This leads to significant time consumption.

Facing the time-consuming problem of the SMS method, we have a key insight that the specular paths tend to be similar among the neighboring shading points. Therefore, we propose to reuse the specular paths spatially. These specular paths can serve as an efficient initialization for the specular manifold walk. This approach results in a significantly reduced average number of Newton iterations required for each shading point. More specifically, we generate some specular path samples with a low sample rate, and then reuse these specular path samples as the initialization for manifold walk among neighboring shading points. Furthermore, we extend the reuse of specular paths from the spatial domain to the temporal domain, accommodating scenes with dynamic lighting or deformation of specular geometry. Our method can provide better quality with equal time. In summary, our main contributions are the following:

- a specular path reuse approach in the spatial domain for efficient caustics rendering,
- a specular path reuse approach in the temporal domain for dynamic light sources and deformation of specular geometry.

2. Related Work

2.1. Caustics Rendering

Over the years, numerous research groups have investigated efficient sampling techniques for challenging specular paths to improve caustic rendering.

Bidirectional Path Tracing [Vea97] employs diverse sampling strategies to connect vertices on subpaths originating from the camera and light source. However, despite the various sampling strategies, this method struggles to handle specular-diffuse-specular (SDS) paths. The limitations stem from the presence of specular and near-specular vertices. Connecting such vertices while satisfying their half-vector constraints is highly improbable.

Metropolis Light Transport [VG97] algorithm introduces the concepts of Markov Chain Monte Carlo (MCMC) methods to rendering. Based on Markov chains, MCMC techniques such as

Hachisuka et al. [HKD14], Šik et al. [SOHK16] and Bitterli et al. [BJNJ17], offer an elegant procedure for exploring the path space by leveraging information from previous iterations. In each iteration, mutations are applied to alter the path structure, and perturbations are used to make small adjustments to vertex positions. However, a drawback of using Markov Chains in this context is the lack of temporal coherence. The Manifold Exploration framework [JM12] incorporates mutation strategies in the MLT context. This framework efficiently renders indirect caustics by employing an iterative process similar to Newton's method, which considers the light path's differential geometry.

Manifold Next Event Estimation [HDF15] introduces an extension of the Next Event Estimation technique by incorporating Manifold Exploration. This extension enabled the connection of surfaces to light sources through a sequence of refractive events, but it can at most find a single specular path due to the simple initialization. This method is effective for rendering simple smooth geometry but encounters challenges when dealing with complex geometry where multiple solutions may exist. MNEE was later extended to bidirectional path tracing by Speierer et al. [SSJ18].

Specular Manifold Sampling [ZGJ20] can find multiple valid specular paths. However, finding a specular path and calculating the corresponding probabilities is time-consuming. We focus on addressing the efficiency issues of this method. In Section 3, we provide a detailed description of the SMS method that is the foundation of our work. Recently, Jhang et al. [JC22] proposed Specular Manifold Bisection Sampling (SMBS), which introduces a large mutation strategy to avoid the issue of converging to a local minimum in the SMS method.

Specular Next Event Estimation [LZHJ20] provides analytic expressions for predicting the total radiance resulting from a single reflective or refractive triangle with a microfacet BSDF. Building on this closed-form solution of a slope-space integral, SNEE achieves efficient caustics rendering, supporting surfaces ranging from ideally specular to near-specular and even highly rough materials. The main limitation of this method is its restriction to specular chains with only a single vertex.

Path Cuts [WHY20] can find multi-bounce pure specular paths from a point light to a pinhole camera. This work was later extended to handle caustic effects [LWT*22] by using the extended specular paths to estimate the incident radiance distribution and guide path tracing. This approach can handle multi-bounce reflective caustics, but it is not accurate enough for near-specular reflectors with very small roughness and does not support refractive surfaces.

2.2. Path Reusing

The path reuse approach is based on the insight that paths from one pixel usually also contribute to its neighboring pixels, and independently sampled paths can be reused after a low-cost shift-mapping operation.

Path Reusing [BSH02] amortizes the cost of path tracing over several screen pixels. As a cost-effective method to generate new path samples for a pixel, Bekaert et al. [BSH02] suggested connecting the primary hit point of the pixel's path to the secondary

hit points of other path samples when they are mutually visible. Gradient-domain path reusing [BPE17] generalizes path reusing by employing the tools for shifting paths that were designed in the context of gradient-domain rendering and improved its efficiency for scenes with highly-glossy and specular light transport.

Reservoir-based Spatiotemporal Importance Resampling (ReSTIR) [BWP*20] employs chained reservoir resampling to share samples across pixels and frames. It improves each pixel's direct light sampling PDF by reusing statistics from temporal and spatial neighbors. Recent work extends ReSTIR to world-space sample reuse [Boi21] and longer paths [OLK*21] for global illumination. Lin et al. [LKB*22] introduced a generalized RIS theory (GRIS) and reformulated ReSTIR's spatiotemporal reuse to long paths and complex specular transport.

Unlike the path reuse approach discussed above, our method reuses admissible specular paths as initial guesses for manifold walk [JM12] to find valid specular paths faster.

3. Background

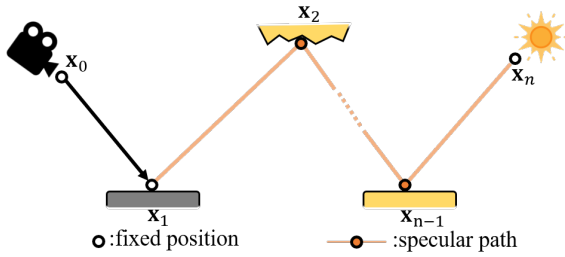


Figure 2: Example specular path between the camera and an emitter position.

Figure 2 shows an example of a complex path containing a chain of specular vertices between two non-specular endpoints \mathbf{x}_1 and \mathbf{x}_n . For example, \mathbf{x}_1 could be a shading point on a diffuse surface and \mathbf{x}_n a position on a light source. The caustics appear on \mathbf{x}_1 , while these specular vertices can occur when light reflects or refracts from a surface, such as a mirror or a glass object. At specular interfaces, light must adhere to the laws of reflection or refraction (Snell's law), significantly reducing the probability of sampling an admissible specular path that connects the endpoints \mathbf{x}_1 and \mathbf{x}_n .

Specular Manifold Sampling (SMS) [ZGJ20] tackles the challenge of sampling such specular paths, including cases when reflection or refraction occurs on a specular surface. It achieves this by performing manifold walk [JM12] involving Newton iteration and reprojection steps to identify solutions that satisfy the constraint function. For high-frequency normal- or displacement-mapped geometry, multiple solutions may satisfy the specular constraints.

The MNEE method [HDF15] performs manifold walks using a fixed initialization and only finds one valid path at most. By randomly sampling the initial guess from a probability distribution $p(\bar{\mathbf{x}}_0)$, the SMS method has the ability to find multiple solutions. The initial guess method is to sample positions uniformly on specular surfaces. To implement this sampling strategy, two uniformly

distributed random numbers $\xi = (\xi_1, \xi_2) \in [0, 1)^2 =: \mathcal{U}$ are taken as input and then warped to the desired distribution.

The obtained sample is utilized as the starting point for the manifold walk algorithm, then perform Newton iteration to find a solution corresponding to a valid specular path. In addition to determining the solutions, it is crucial to calculate the discrete probability, denoted as p_k , of identifying a specific path vertex $\mathbf{x}_2^{(k)}$.

When analyzing the convergence behavior of the manifold walks on the *primary sample space* \mathcal{U} , Zeltner et al. [ZGJ20] observe multiple *basins of convergence* $\mathcal{B}_k \subseteq \mathcal{U}$, each containing a point $\xi^{(k)}$ identified with a corresponding solution vertex $\mathbf{x}_2^{(k)}$. They propose that the probability p_k can be determined as the area of the corresponding convergence basin on \mathcal{U} :

$$p_k = \int_{\mathcal{U}} \mathbb{1}_{\mathcal{B}_k}(\xi) d\xi. \quad (1)$$

Zeltner et al. [ZGJ20] introduce two algorithms to estimate this integral: an unbiased estimate of the integral named unbiased SMS and a biased but consistent estimate named biased SMS. While the biased SMS technique offers lower running time and variance than the unbiased SMS, it still requires multiple iterations to search for specular paths.

The biased SMS algorithm uses a fixed number of M samples for Bernoulli trial iterations. These M samples are grouped into a set of unique solutions $\mathbf{x}_2^{(l)} (l = 1, \dots, L)$. The biased estimate of the integral reciprocal is obtained by calculating the relative frequency of occurrence n_l for each solution. The total throughput at shading point \mathbf{x}_1 is then estimated using the biased estimator:

$$\frac{1}{M} \sum_{l=1}^L n_l \frac{f(\mathbf{x}_2^{(l)})}{p(\mathbf{x}_2^{(l)})} \approx \frac{1}{M} \sum_{l=1}^L n_l f(\mathbf{x}_2^{(l)}) \frac{M}{n_l} = \sum_{l=1}^L f(\mathbf{x}_2^{(l)}). \quad (2)$$

The biased SMS algorithm introduces bias compared to the original unbiased method, which results in energy loss visible as darkening in the caustic regions of the final image. As M increases, the energy lost decreases, and the biased SMS algorithm converges to the actual value when M approaches infinity.

Our method for reusing specular paths in both spatial and temporal domains is specifically designed for the biased SMS algorithm, which cannot be directly applied to the unbiased algorithm due to the challenge of determining the integral value of $\langle 1/p_k \rangle$ after reusing. In contrast, the biased algorithm circumvents this issue and allows for effective path reuse.

In order to efficiently find valid solutions for high-frequency normal-mapped surfaces, Zeltner et al. [ZGJ20] proposed a *two-stage* sampling approach. In the first stage, a smooth manifold walk is performed using a Gaussian approximation of the specified normal map, which brings the algorithm closer to the final solution. In the second stage, on the actual bumpy surface, the manifold walk is started from the position obtained in the first stage to reach a valid solution.

4. Our approach

4.1. Problem analysis and motivation

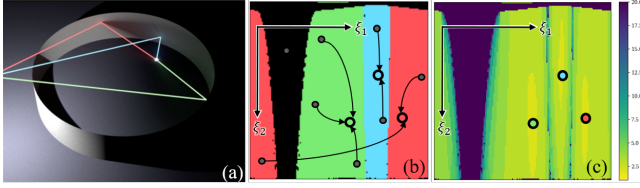


Figure 3: (a): The three solution paths at a particular shading point in the Ring scene. (b): The basins of convergence (in primary sample space) corresponding to those solution paths. (c): The visualization of the number of Newton iterations required for different initial points.

Zeltner et al. [ZGJ20] treat both the smooth surface and the rough surfaces in a universal way, by finding admissible specular paths that satisfy the constrain functions. The core of their method is sampling two random numbers uniformly in the primary sample space and mapping them to positions on the specular geometry. These mapped points are the starting points for Newton iteration to generate admissible specular paths.

As shown in Figure 3b, random numbers in different convergence basins lead to different solutions, while those in the same convergence basin converge to the same solution. The number of Newton iterations may vary depending on the distance between the initial point and the final solution (Figure 3c). Generally, the closer the initial point is to the solution, the fewer iterations are required. The number of Newton iterations performed during the manifold walk significantly affects the running time of the SMS algorithm. One key question is how to set the initial random number, which corresponds to the starting point on specular geometry, close to the final solution.

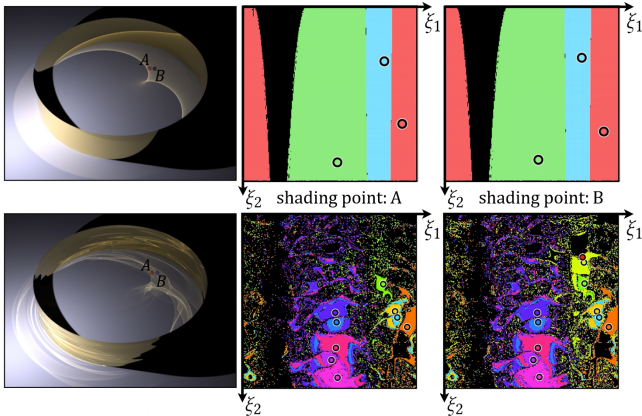


Figure 4: The convergence basins of two neighboring shading points in the Ring scene, without and with a normal map. The shape of their respective convergence basins is similar, and their solutions are close. Top: For the Ring scene without a normal map. Bottom: For the Ring scene with a normal map.

Our key insight is that neighboring shading points tend to have

similar solutions in the primary sample space. Figure 4 shows the converged basins of two neighboring shading points in the Ring scene, without and with a normal map. For the simple case, the shapes of their convergence basins are similar, and their solutions in the primary sample space are close. This similarity holds even when the specular geometry includes a normal map. Additionally, it remains valid for shading points that do not have solutions, indicated by completely black convergence basins. It is easy to associate that points similar in the primary sample space are also similar in the world space.

Building on this observation, we propose to reuse admissible specular paths as the initial guess for manifold walk in the spatial domain. In regions without caustics where no solution exists, we utilize the spatial neighborhood information to determine whether to perform manifold walk. We extend this reuse strategy from the spatial domain to the temporal domain for scenes with dynamic lighting or deformation of specular geometry.

4.2. Solution Overview

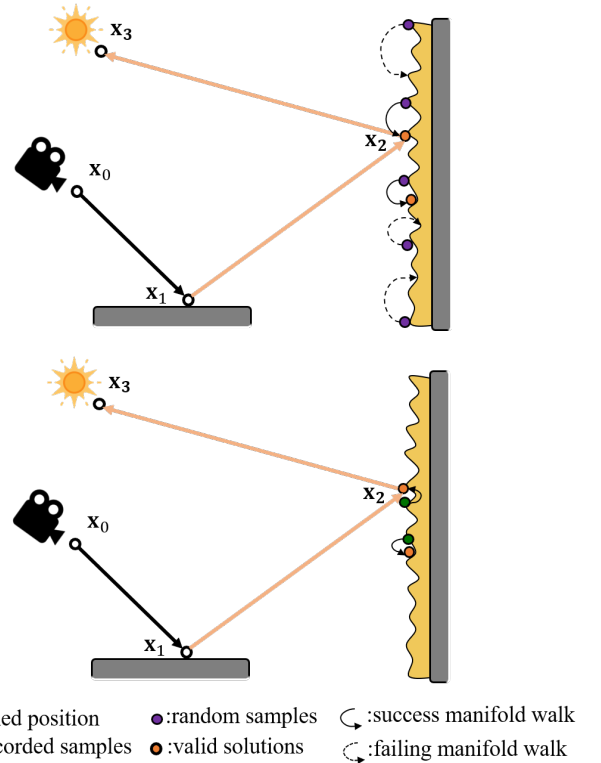


Figure 5: The process of finding specular paths using specular manifold in the SMS [ZGJ20] method (top) and our method (bottom). The SMS method obtains the starting points by sampling uniformly over the specular geometry, while our method starts from the recorded samples. These recorded specular path samples are obtained through spatial and temporal reuse. Compared with the SMS method, our method uses the recorded samples as the starting points of manifold walk, leading to a higher manifold walk success rate and fewer Newton iterations.

Our basic idea is to reuse the admissible specular paths as initial guesses for manifold walk to construct valid specular paths faster. Figure 5 shows the different ways of choosing the starting points between our method and the SMS [ZGJ20] method. Our method consists of the following steps:

- In the first phase, we generate admissible specular paths with a lower sample rate using the original biased Specular Manifold Sampling (SMS) algorithm (Sec. 4.3).
- In the second phase, we obtain information from the admissible specular paths generated in the first stage, including the number of specular paths and their corresponding points on the specular geometry. We share this information in the spatial domain and apply it to all newly generated shading points (Sec. 4.4).
- For scenes with dynamic lighting or deformation of specular geometry, we also employ this reusing process in the temporal domain by utilizing information obtained from the previous frame (Sec. 4.5).

Our method is suitable for rendering caustics from point and small area light sources and supports both reflection and refraction. For glossy scenes with some roughness, our method also brings enhancements. Furthermore, it can be extended to handle multiple-bounce scenes.

4.3. Specular path sample generation

Our goal is to reuse the specular paths among nearby shading points. To achieve this, we generate specular path samples following the biased SMS algorithm proposed by Zeltner et al. [ZGJ20], but using a less costly number of Bernoulli trials M and a lower sampling rate.

When using multiple samples per pixel (SPP) for path sample generation, it is likely to generate *duplicate* (not exactly the same) specular paths. To address this issue, we employ a *de-duplication* strategy when merging path samples generated by multiple SPP.

Note that this step is quite efficient since the specular path samples are generated with much less budget than the rendering.

4.4. Spatial reuse of specular path samples

According to our observation, the specular paths are similar at neighboring shading points, and this observation holds within the same pixel and between adjacent pixels. Thus with the information about the specular path samples generated within each pixel, we reuse them in the spatial domain.

Considering that the pixels are rendered in Morton code order, the size of the spatial reuse range can be set to $2^n \times 2^n$, where $n \geq 0$. Pixels in the same reuse range share the specular path sample information of each other, but those shared paths are de-duplicated based on their respective pixel's shading point information during merging.

After the spatial merging process at each pixel, our method reuses the obtained information. First, our method utilizes the number of recorded specular paths to set different numbers of Bernoulli trials for different pixels of the image rather than a fixed value. Second, our method utilizes the recorded specular vertices as the

Table 1: Comparison of the average number of Newton iterations for all test scenes. The SMS employs a two-stage sampling approach in the first three scenes with normal map, while it does not utilize the two-stage sampling in the Slab scene without a normal map. Our approach only utilizes the second manifold walk during the reuse phase. Our method shows the average number of Newton iterations in the reuse phase. The average number of Newton iterations calculation includes both success and failing manifold walks.

Scene	SMS		Ours
	first stage	second stage	second stage
Ring (normal-mapped)	16.7	8.7	2.5
Reflective plane	10.5	13.5	1.9
Refractive sphere	11.2	6.5	2.9
Double-refractive slab	/	18.4	4.1

starting points for the manifold walk, and then performs Newton iterations to find valid specular paths that satisfy the constraint function. After traversing all recorded specular vertices, we collect the unique specular paths found based on them to calculate the total throughput of the current shading point (according to equation 2).

Discussion. Our method improves the efficiency of the SMS method by optimizing in two aspects. On the one hand, Zelter et al. [ZGJ20] apply the SMS algorithm to the entire image, aiming to discover potential specular paths. However, this approach often involves numerous unsuccessful attempts, leading to a high computational cost. Our method introduces a different number of Bernoulli trials at each pixel. This adaptive value allows us to avoid unnecessary exploration in regions without caustics, significantly improving efficiency. On the other hand, our method reduces the number of Newton iterations by reusing the recorded specular vertex as the initial points when dealing with high-frequency geometry. The average number of Newton iterations of our method and the original biased SMS [ZGJ20] is shown in Table 1. Obviously, our method significantly reduces the number of iterations.

For glossy scenes with some roughness, one way to improve the efficiency of the SMS method is by increasing the spatial reuse size. By obtaining more path samples through a larger reuse size, we can better explore the space of specular paths and increase the chances of finding valid solutions. This can lead to improved efficiency in rendering glossy scenes with rough surfaces.

4.5. Temporal reuse of specular path samples

Our reuse process extends to the temporal domain for scenes with dynamic lighting or deformation of specular geometry. We leverage information from the previous frame by merging the recorded specular path samples pixel by pixel with the samples obtained in the first phase of the current frame, using the same *de-duplication* strategy.

While rendering the current frame, we utilize the merged information to enhance our results. Moreover, we can continue to apply spatial reuse techniques, taking advantage of the merged informa-

tion to further improve rendering quality. Figure 6 illustrates the data flow of the specular path samples generated in the first phase.

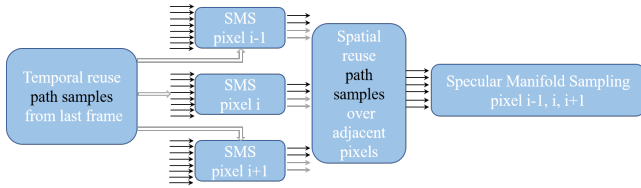


Figure 6: We store the path samples from the previous frame and combine them with the samples generated during the rendering of the first phase of the current frame.

4.6. Implementation

Our approach uses the angle difference constraint function to perform Newton iterations. In addition, we adopt the *two-stage* manifold walk to generate valid specular paths (Sec.4.3). And we only utilize the second manifold walk during the reuse phase (Sec.4.4; Sec.4.5).

When rendering each pixel, the biased SMS algorithm may be invoked at different path depths to search for multiple specular paths. When a specular path is found, we record the shading point on the diffuse geometry with its corresponding path depth and the vertex on the specular geometry. For two-bounce caustics, we record the vertex on the first specular geometry. We record information about these specular path samples in the screen space. We also record the number of specular paths found at different path depths of the pixel.

To solve the problem of duplicate solutions in the specular path sample generation phase, we take the following approach: If the vertices on the specular geometry of two specular paths are close, we record only one of them. Our first samples per pixel (SPP) forces light to be emitted from the center of the pixel. We record the position of the shading point. Then, we connect these two specular vertices to the corresponding shading point in the first SPP and calculate their directional difference. If the difference is less than a threshold, we consider it the same. It might be helpful to consider the geometry normal information of the specular vertices in the future.

5. Results

We implemented our method in Mitsuba 2 [NDVZJ19], building upon the implementation of the SMS method by Zeltner et al. [ZGJ20]. All time measurements in this section were recorded on a 3.20GHz Intel i9 CPU (24 cores) with 64GB of RAM. We compare against path tracing (PT), specular manifold sampling (SMS) [ZGJ20] and specular next event estimation (SNEE) [LZHJ20] with equal time, and a converged PT as reference.

Ring scene. In Figure 7, we present the results of our method for the entire image and compare the effectiveness with brute-force path tracing and the state-of-the-art method SMS (biased) on the Ring scene with a normal map in equal time. Our method employs 2 SPP in the path sample generation phase, with M set to the

same value as in the biased SMS method. The results demonstrate that our method surpasses the biased SMS approach, exhibiting improved brightness that closely matches the reference image in regions with intricate caustics. Moreover, as the spatial reuse range increases, the quality of the results continues to improve.

Plane and Sphere scenes. In Figure 8, we compare the effectiveness of our method, path tracing, and biased SMS with equal time. We do not compare our method with SNEE [LZHJ20] in this context as it is designed for handling actual displacement geometry. Our method can achieve more SPP compared to the biased SMS method. Moreover, our method effectively recovers lost energy and produces a caustic brightness resembling the reference image.

Double refractive slab scene. Our method can handle long chains with multiple specular interactions. For two-bounce caustics, as shown in Figure 1, we use the specular path samples generated in the first phase to determine the initial point on the first interface (bumped glass), and the initial point on the second interface (planar glass) by tracing a ray through the first vertex. We compare our result with path tracing, SNEE, biased SMS with equal time. Since SNEE cannot handle multiple intermediate specular interactions, we only observe caustics in the area between the two glass panes, while the rest of the image performs similarly to the PT method. In contrast, our method still produces good results.

Spatial reuse size. In Figure 9, we show the performance of our method under different spatial reuse size. As the range of spatial reuse increases, the rendering time increases while the image quality improves. However, this improvement is not linear, and a larger size of reuse becomes unnecessary at a certain point. In our experiments, we observe that different scenes require selecting the appropriate spatial reuse size based on their specific characteristics and also related to the number of SPP used in the path sample generation phase.

Roughness. In Figure 10, we compare our method, PT, biased SMS on the rough Plane scene in equal time. As roughness increases, PT becomes more capable of directly connecting to the light source, whereas the SMS method loses efficiency. By increasing the size of spatial reuse, our approach can achieve better results compared to the SMS method.

Moving light source. In addition to spatial reuse, our approach can be extended to scenes with dynamic lighting, allowing us to utilize information from the temporal domain. In Figure 11, we compare the results of our method with and without temporal reuse in equal time. The result with temporal reuse utilizes the path sample information recorded before the rotation of the light source in the previous frame. The results demonstrate that temporal reuse can improve the rendering quality of complex caustic regions.

Deformation of specular geometry. Our method can also be used for movement and deformation of specular geometry. In Figure 12, We show the results of deforming the specular geometry. We perform a Gaussian filter operation on the normal map to deform the specular geometry. The result of temporal reuse utilizes the path sample information generated by the more complex specular geometry of the previous frame. The results demonstrate that temporal reuse is also effective in scenes involving specular geometry defor-

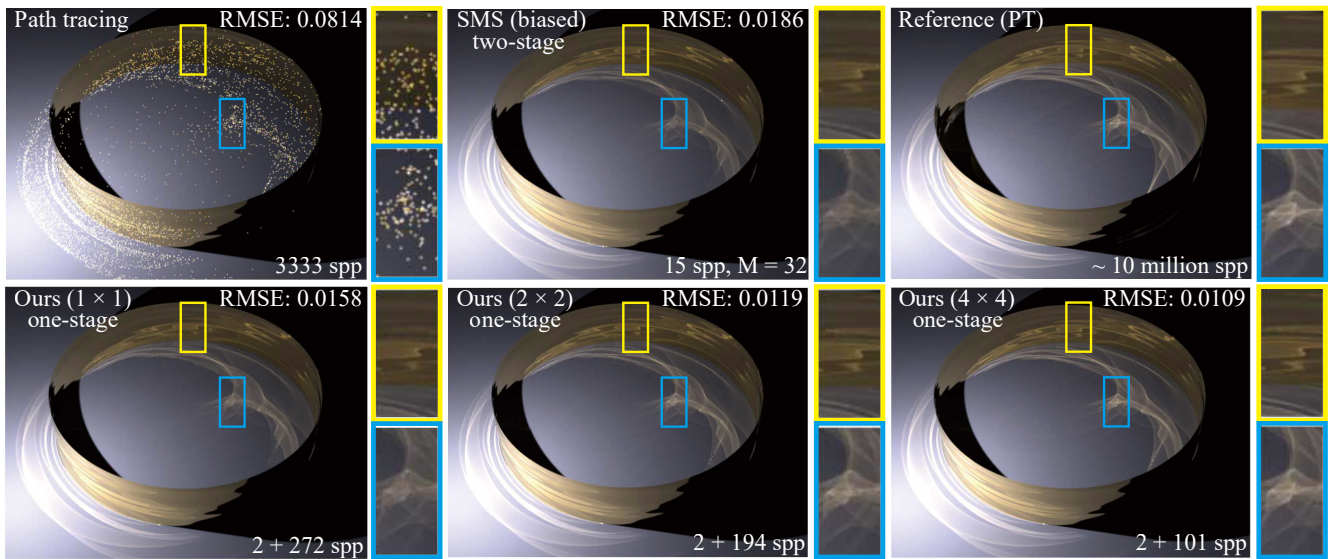


Figure 7: Equal-time comparison (1 minute) among our method, path tracing, biased SMS [ZGJ20] on the ring scene with a normal map. We report the number of SPP computed by each method and RMSE compared to the reference image. Our method employs 2 SPP in the path sample generation phase, with M set to the same value as in the biased SMS method. The results of different spatial reuse size are also shown.

mation and can further recover the lost caustic energy to improve rendering quality.

Difference map. In Figure 13, we provide visualizations to illustrate the differences between the results and the reference image. The results show that spatial reuse reduces the caustic energy loss in all image regions, while temporal reuse can further improve the complex caustic regions.

6. Discussion and Limitations

We discuss the main limitations of our method as follows.

Selection of parameters. Our approach involves the choice of some parameters, such as the number of SPP and Bernoulli trials M in the path sample generation phase, the spatial reuse size, with or without temporal reuse. As illustrated in Figure 9, an excessively large spatial reuse range results in marginal improvements in image quality while significantly increasing rendering time. Our approach requires finding the optimal balance in the specific settings of these parameters.

Biasedness and Consistency. Our method is biased and consistent. We compute the probability of the resulting specular paths after reuse in the same way as the biased SMS method. Therefore, our method is also biased. Our approach cannot be applied to unbiased SMS algorithms due to the challenging nature of calculating the probability of specular paths generated after reuse in an unbiased manner. To ensure the efficiency of our method, we set the number of Bernoulli trials using only the specular path samples generated at a low sample rate. We can increase the sampling rate and the number of Bernoulli trials in the first phase or increase the spatial reuse range and the number of temporal reuse frames in the second phase of our method to increase the number of Bernoulli trials, thus converging to the actual value.

Limited roughness. As shown in Figure 10, our method exhibits reduced effectiveness when dealing with higher roughness. The higher roughness leads to a significant discrepancy in the *offset normal* sampled during the *offset manifold walk* [JM12].

Large light sources. Our approach is well-suited for rendering caustics from point and small area light sources. However, its effectiveness diminishes when dealing with large light sources such as environment maps or large area light sources. This is due to the significant differences in specular paths between neighboring shading points in such cases, causing a reduced success rate in manifold walk for reused paths.

Storage. Our method requires storing the generated specular paths, which incurs additional storage requirements. The storage size depends on the image resolution, and the number of SPP and Bernoulli trials M used during the specular path sample generation phase. Our experiments find that the total size of the extra storage for each scene is at most 200MB.

7. Conclusion

We propose an efficient caustics rendering method via spatial and temporal path reusing. Our basic idea is to reuse the specular paths generated by the specular manifold sampling method with a low sample rate, then reuse these specular paths as starting points for manifold walk to construct valid specular paths faster among neighboring shading points. Furthermore, we extend this reuse strategy from spatial domain to temporal domain, for scenes with dynamic lighting or deformation of specular geometry. Our method outperforms current state-of-the-art methods and can handle multiple bounces of light and a variety of scenes.

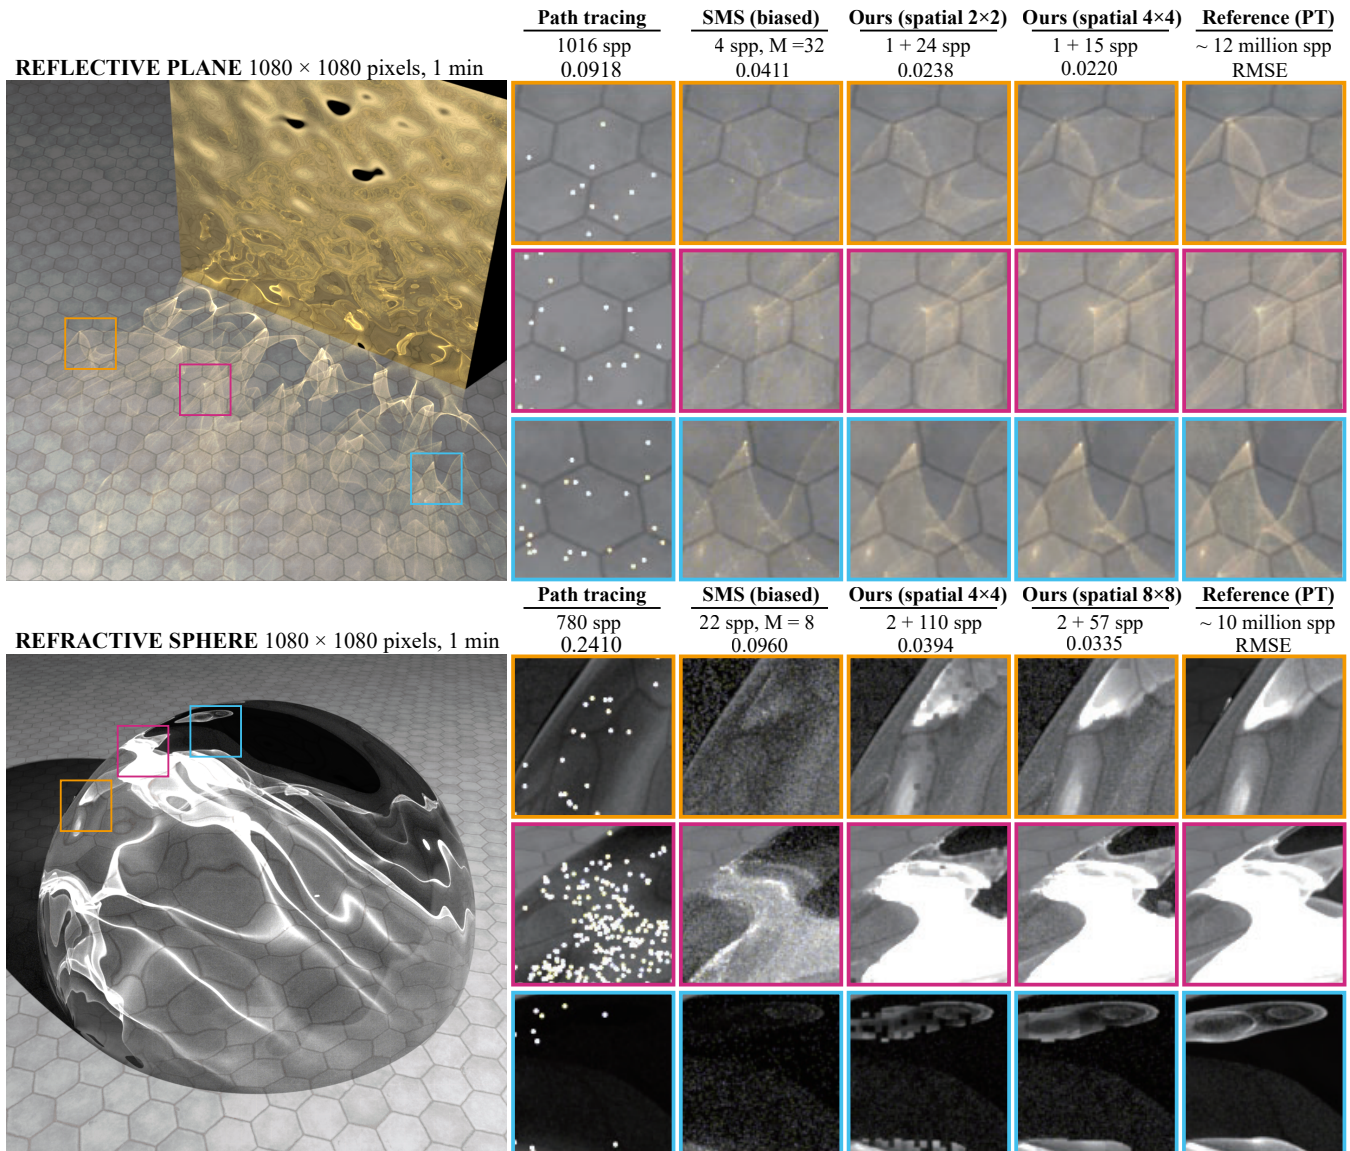


Figure 8: Equal-time comparison (1 minute) among our method, path tracing, biased SMS [ZGJ20] on the single-bounce reflective and refractive scene with challenging caustics due to normal-mapped surfaces. Our method employs 1 SPP for the Plane scene and 2 SPP for the Sphere scene in the path sample generation phase, with M set to the same value as in the respective biased SMS method. We report the number of SPP computed by each method and RMSE compared to the reference image.

Acknowledgments

We thank the reviewers for the valuable comments. This work has been partially supported by the National Key R&D Program of China under grant No. 2022ZD0116305, National Natural Science Foundation of China under grant No. 62272275 and No. 62172220.

References

- [BJNJ17] BITTERLI B., JAKOB W., NOVÁK J., JAROSZ W.: Reversible jump metropolis light transport using inverse mappings. *ACM Trans. Graph.* 37, 1 (oct 2017). URL: <https://doi.org/10.1145/3132704>, doi:10.1145/3132704. 2
- [Boi21] BOISSÉ G.: World-space spatiotemporal reservoir reuse for ray-traced global illumination. In *SIGGRAPH Asia 2021 Technical Communications* (New York, NY, USA, 2021), SA '21 Technical Communications, Association for Computing Machinery. URL: <https://doi.org/10.1145/3478512.3488613>, doi:10.1145/3478512.3488613. 3
- [BPE17] BAUSZAT P., PETITJEAN V., EISEMANN E.: Gradient-domain path reusing. *ACM Trans. Graph. (Proc. of SIGGRAPH Asia)* 36, 6 (November 2017). URL: <http://graphics.tudelft.nl/Publications-new/2017/BPE17>. 3
- [BSH02] BEKAERT P., SBERT M., HALTON J.: Accelerating path trac-

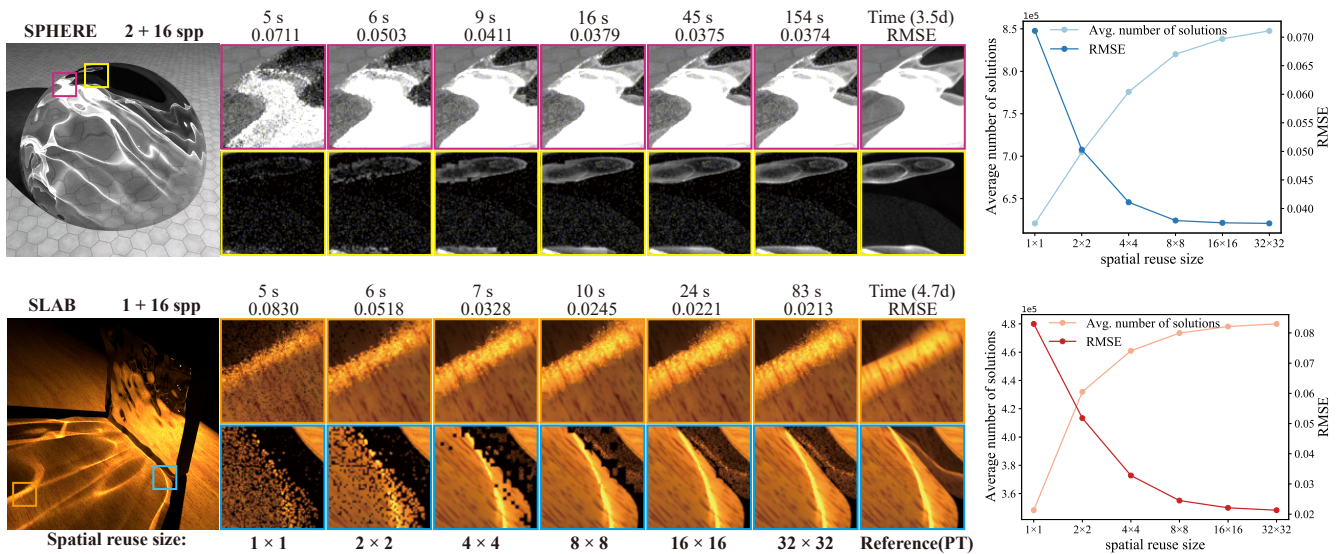


Figure 9: Analyzing the performance of our approach on Sphere and Slab scene with different spatial reuse sizes. Top: Our method employs 2 SPP in the path sample generation phase and 16 SPP in the spatial reuse phase for the Sphere scene. Bottom: Our method employs 1 SPP in the path sample generation phase and 16 SPP in the spatial reuse phase for the Slab scene. The two dashes in the plots depict how the number of admissible specular paths found and the image quality (RMSE) change as the spatial reuse size increases.

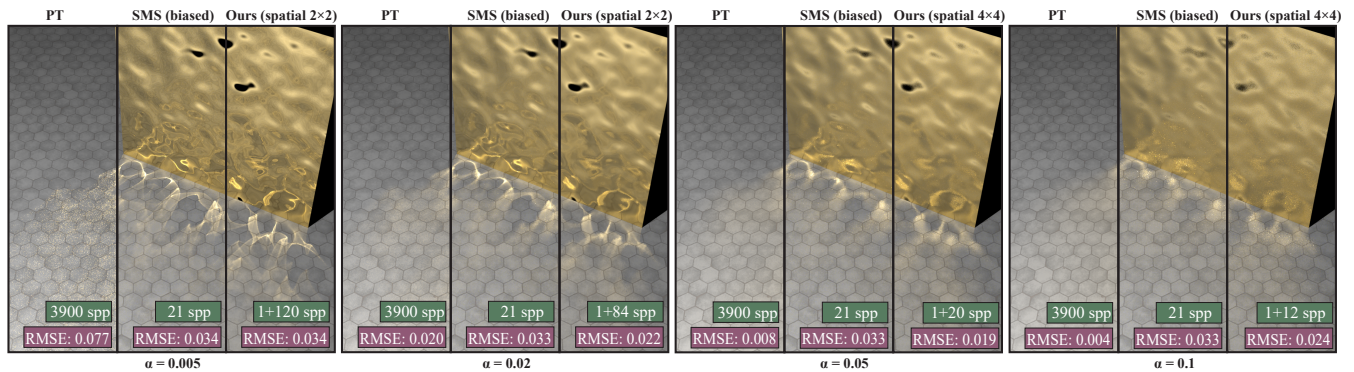


Figure 10: Equal-time comparison (5 minutes) among our method, PT, biased SMS [ZGJ20] on the rough reflective Plane scene. Insets show the number of SPP and RMSE compared to the reference image. As the roughness increases, the path tracing method becomes more capable of directly connecting to the light source, while the SMS method loses efficiency. By increasing the size of spatial reuse in our approach, we are able to achieve better results compared to the biased SMS method.

- ing by re-using paths. In *Proceedings of the 13th Eurographics Workshop on Rendering* (Goslar, DEU, 2002), EGRW '02, Eurographics Association, p. 125–134. 2
- [BWP*20] BITTERLI B., WYMAN C., PHARR M., SHIRLEY P., LEFOHN A., JAROSZ W.: 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/gg8xc7.3](https://doi.org/10/gg8xc7.3)
- [HDF15] HANIKA J., DROSKE M., FASCIONE L.: Manifold Next Event Estimation. *Computer Graphics Forum* (2015). 2, 3
- [HKD14] HACHISUKA T., KAPLANYAN A. S., DACHSBACHER C.: Multiplexed metropolis light transport. URL: <https://doi.org/10.1145/2601097.2601138>, doi:10.1145/2601097.2601138. 2
- [JC22] JHANG J.-W., CHANG C.-F.: Specular Manifold Bisection Sampling for Caustics Rendering. *Computer Graphics Forum* (2022). doi: [10.1111/cgf.14673.2](https://doi.org/10.1111/cgf.14673.2)
- [JM12] JAKOB W., MARSCHNER S.: 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. 2, 3, 7
- [Kaj86] KAJIYA J. T.: The rendering equation. *SIGGRAPH Comput. Graph.* 20, 4 (aug 1986), 143–150. URL: <https://doi.org/10.1145/15886.15902>, doi:10.1145/15886.15902. 2
- [LKB*22] LIN D., KETTUNEN M., BITTERLI B., PANTALEONI J., YUKSEL C., WYMAN C.: Generalized resampled importance sampling: foundations of restr. *ACM Transactions on Graphics (TOG)* 41, 4 (2022), 1–23. 3
- [LWT*22] LI H., WANG B., TU C., XU K., HOLZSCHUCH N., YAN L.-Q.: Unbiased caustics rendering guided by representative specular paths. In *SIGGRAPH Asia 2022 Conference Papers* (New York,




Figure 11: Equal-time comparison (1 minute) of our method w/o (left) and w/ (right) temporal reuse on the moving light source scene. We employ 2 SPP in the path sample generation phase and spatial reuse size is 4×4 . The SPP for rendering and RMSE compared to the reference image are also provided.




Figure 12: Equal-time comparison (1 minute) of our method w/o (left) and w/ (right) temporal reuse on the specular geometry deformation scene. We employ 1 SPP in the path sample generation phase, and spatial reuse size is 2×2 . Top: The result of using a more complex normal map. Bottom: The result of using a simpler normal map. The SPP for rendering and RMSE compared to the reference image are also provided.

NY, USA, 2022), SA '22, Association for Computing Machinery. URL: <https://doi.org/10.1145/3550469.3555381>, doi: 10.1145/3550469.3555381. 2

- [LZHJ20] LOUBET G., ZELTNER T., HOLZSCHUCH N., JAKOB W.: Slope-space integrals for specular next event estimation. *Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 39, 6 (dec 2020). 1, 2, 6
- [NDVZJ19] NIMIER-DAVID M., VICINI D., ZELTNER T., JAKOB W.: Mitsuba 2: A retargetable forward and inverse renderer. *Transactions on Graphics (Proceedings of SIGGRAPH Asia)* 38, 6 (Dec. 2019). doi: 10.1145/3355089.3356498. 6
- [OLK*21] OUYANG Y., LIU S., KETTUNEN M., PHARR M., PANTALEONI J.: Restir gi: Path resampling for real-time path tracing. In *Computer Graphics Forum* (2021), vol. 40, Wiley Online Library, pp. 17–29. 3
- [SOHK16] SIK M., OTSU H., HACHISUKA T., KRIVÁNEK J.: Robust light transport simulation via metropolised bidirectional estimators. *ACM Trans. Graph.* 35, 6 (dec 2016). URL: <https://doi.org/10.1145/2980179.2982411>, doi:10.1145/2980179.2982411. 2

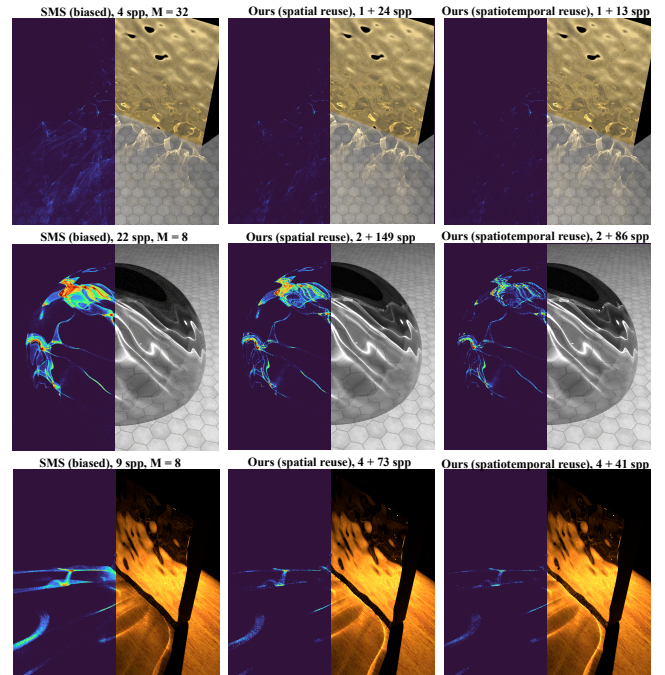


Figure 13: Equal-time comparison (1 minute) among biased SMS, our method (spatial reuse), and our method (spatiotemporal reuse). Overlays show the pixel-wise squared error compared to the reference image.

- [org/10.1145/2980179.2982411](https://doi.org/10.1145/2980179.2982411), doi:10.1145/2980179.2982411. 2
- [SSJ18] SEBASTIEN SPEIERER CHRISTOPHE HERY R. V., JAKOB W.: Caustic connection strategies for bidirectional path tracing. Pixar Technical Memo 18-01, 2018. 2
- [Vea97] VEACH E.: *Robust Monte Carlo Methods for Light Transport Simulation*. PhD thesis, Stanford University, 1997. 2
- [VG97] VEACH E., GUIBAS L. J.: Metropolis light transport. In *Computer Graphics (ACM SIGGRAPH '97 Proceedings)* (1997), vol. 31, pp. 65–76. 2
- [WHY20] WANG B., HAŞAN M., YAN L.-Q.: Path cuts: Efficient rendering of pure specular light transport. *ACM Trans. Graph.* 39, 6 (Nov. 2020). 2
- [ZGJ20] ZELTNER T., GEORGIEV I., JAKOB W.: Specular manifold sampling for rendering high-frequency caustics and glints. *ACM Trans. Graph.* 39, 4 (jul 2020). 1, 2, 3, 4, 5, 6, 7, 8, 9