

# Supplementary Material: Unbiased Caustics Rendering Guided by Representative Specular Paths

He Li  
Shandong University  
China  
he.li@mail.sdu.edu.cn

Beibei Wang\*  
Nankai University  
China  
beibei.wang@nankai.edu.cn

Changhe Tu\*  
Shandong University  
China  
chtu@sdu.edu.cn

Kun Xu  
Tsinghua University  
China  
xukun@tsinghua.edu.cn

Nicolas Holzschuch  
University Grenoble Alpes, Inria,  
CNRS, Grenoble INP, LJK  
France  
nicolas.holzschuch@inria.fr

Ling-Qi Yan  
University of California, Santa  
Barbara  
USA  
lingqi@cs.ucsb.edu

## CCS CONCEPTS

• Computing methodologies → Ray tracing.

## 1 SUPPLEMENTARY MATERIAL

### 1.1 Finding the reflector point within a triangle

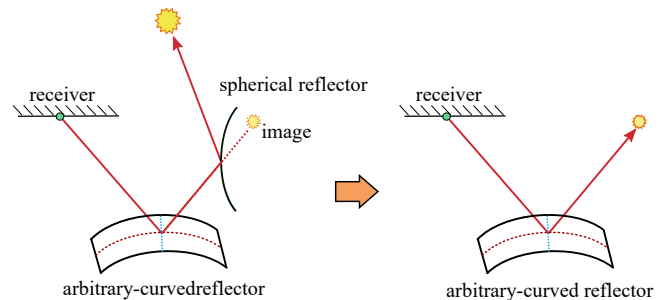
For each triangle of a leaf path cut, we use the Newton solver to find the reflector point within the triangle to construct a valid path. During the Newton iteration, if a point locates outside the triangle, we move the point back to the nearest edge or corner of the triangle and continue the Newton iteration until finding a valid path or reaching the maximum step count. During this process, we record the point with the smallest error. Finally, if no valid specular path within the path cut is found, we use the recorded points to form a path, which is ensured to be inside the glossy path region. Though such a path might not be a pure specular path, it is still representative of the glossy reflections in the path cut.

### 1.2 Blurring of the visualization

In Fig. 6 (main paper), we compare our sampling distribution approximated by an SG and the target incident radiance distribution. Our distribution is able to show the distribution shape of the target function, although with some blurring. The blurring is due to several reasons. As a rough approximation, the SG product integral approximation [Xu et al. 2014] would introduce some overblurring, resulting in our “blurred” distribution. Another reason is that the curved reflector would scale the reflection (e.g., a convex mirror would produce a shrunk reflection image), which is not considered in our approximation. Nevertheless, our distribution is able to cover most of the target distribution and enables high-quality importance sampling.

### 1.3 Parallax compensation for multiple-bounce cases

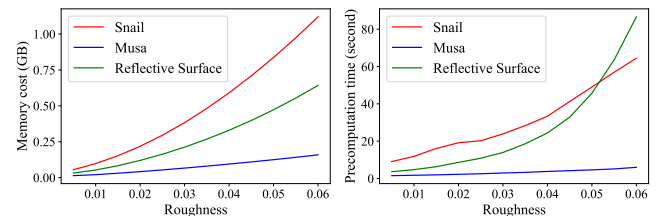
As shown in Fig. 1, we start from the light source, compute the position of the image produced by the reflector (treated as a spherical reflector), and then treat the image as a light source. We repeat this



**Figure 1:** For a spherical reflector, the position of the virtual image can be directly computed [Fitzpatrick 2007]. By treating the image produced by the spherical reflector as a light source, we can reduce multiple-bounce cases (left) to the single-bounce case (right).

process until the last reflector. This reduces multiple-bounce cases to the single-bounce case. Then, we perform the arbitrary-curved parallax compensation for the last reflector.


### 1.4 More limitations of our method



**Figure 2:** Impact of varying roughness on the memory cost and precomputation time. For the Reflective Surface scene, which includes SDS caustics, the threshold  $\epsilon$  of relaxed path cuts is set as 0.1 (set as 0.01 for other scenes).

*Limited roughness.* Our method is suitable for reflectors with low roughness (from about 0.005 to 0.05). For higher or lower roughness,

\*Corresponding authors.



**Figure 3: Equal-time comparison (3 minutes) between our method and path tracing on reflectors with varying roughness. Due to the increasing number of path cuts, the performance of our method degrades with the increasing roughness.**

our method will have performance degradation. For higher roughness, the path cuts method [Wang et al. 2020] becomes slow, due to the large number of valid path cuts and the inefficient pruning of triangles. The number of representative grows fast, resulting in much more memory cost and precomputation time, as shown in Fig. 2. The rendering also becomes slower, because there are more SGs to be sampled. As shown in Fig. 3, when the roughness becomes higher, our method would take fewer samples within the same time, resulting in lower-quality results.

For lower roughness, our approach is less competitive than other methods such as SMS [Zeltner et al. 2020], as the SG approximation and parallax compensation become less accurate, leading to an inaccurate guiding. Precisely, the SG approximation makes the distribution over-blurry when the roughness is extremely low, while the target distribution is very sharp, which reduces the sampling quality. Furthermore, the parallax compensation becomes less accurate for sharply changed incident radiance. We could alleviate this issue with dense cached points, but this could reduce performance and increase memory cost drastically, making our approach less practical.

*Roughness maps.* Our method does not support roughness maps for now. For spatially-varying roughness, we have to use the highest roughness in the range of varying roughness for both path space traversal and SG approximation, which will have a significant impact on performance.

*Large light source.* Our method is designed for point lights and small area lights that are close to point lights. For an area light that cannot be treated as a point light, we have to sample points on the light surface, treat them as separate point lights, and run the algorithm separately. In this way, the performance degrades linearly with the number of the sampled points.

*Refraction.* We do not consider refractive caustics for now, but it's straightforward to extend our method to handle refractive caustics, by following the original path cuts method [Wang et al. 2020]. The main difference is the parallax compensation. Refraction at spherical boundary also produces image, but the method to compute image position [Chu 2001] is different from spherical mirror reflection.

## REFERENCES

- William Chu. 2001. Refraction at Spherical Surfaces. <https://personal.math.ubc.ca/cass/courses/m309-01a/chu/project-notes.htm>.
- Richard Fitzpatrick. 2007. Image Formation by Convex Mirrors. <https://farside.ph.utexas.edu/teaching/316/lectures/node138.html>.
- 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.
- Kun Xu, Yan-Pei Cao, Li-Qian Ma, Zhao Dong, Rui Wang, and Shi-Min Hu. 2014. A practical algorithm for rendering interreflections with all-frequency brdfs. *ACM Transactions on Graphics (TOG)* 33, 1 (2014), 1–16.
- 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.