

Efficient Visibility Reuse for Real-time ReSTIR

Yusuke Tokuyoshi
yusuke.tokuyoshi@amd.com
Advanced Micro Devices, Inc.
Tokyo, Japan

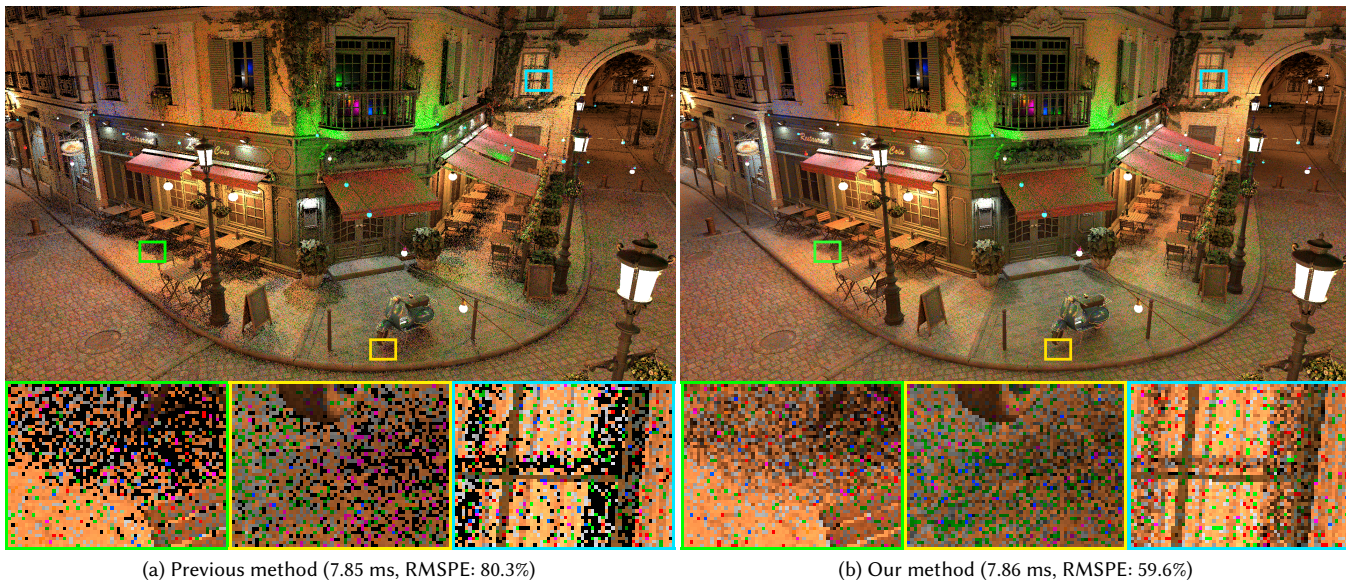


Figure 1: Real-time direct illumination using biased visibility-reuse ReSTIR [Bitterli et al. 2020; Wyman and Pantelev 2021] with and without our shadow variance reduction technique (1920×1080 pixels, two rays per pixel, AMD Radeon™ RX 7900 XTX GPU). Our method reduces the root mean square percentage error (RMSPE) by efficiently reusing spatiotemporal sample visibilities.

ACM Reference Format:

Yusuke Tokuyoshi. 2024. Efficient Visibility Reuse for Real-time ReSTIR. In *Special Interest Group on Computer Graphics and Interactive Techniques Conference Talks (SIGGRAPH Talks '24)*, July 27 - August 01, 2024. ACM, New York, NY, USA, 2 pages. <https://doi.org/10.1145/3641233.3664320>

1 INTRODUCTION

Spatiotemporal reservoir resampling (ReSTIR) [Bitterli et al. 2020] is a powerful sampling technique that is often used for real-time rendering. It samples an important light path by reusing many samples from spatiotemporally neighboring pixels. However, ReSTIR can produce significant noise around shadow edges, because the generated sample distribution is often dissimilar to the illumination integrand around such edges. Since shadow ray visibility is a step function for opaque surfaces, one-sample visibility is insufficient for those shadows. To reduce the shadow noise for ReSTIR, we present a simple and efficient visibility estimation that reuses the visibilities

of spatiotemporal samples. Our method is based on the visibility reuse proposed by Wyman and Pantelev [2021]. While they reused visibilities to reduce shadow rays, we reduce the variance by introducing a variant of *weighted importance sampling* (WIS) [Bekaert et al. 2000]. Although our method can introduce a negligible bias for typical real-time implementations, we also discuss conditions to obtain unbiased results.

2 BACKGROUND

ReSTIR first generates an initial sample for each pixel and then resamples a sample X from spatiotemporally neighboring pixels i according to a weight $w_i \in [0, \infty)$. For the rendering integral, the one-sample estimator is written as

$$\int_{\Omega} f(x)V(x)dx \approx f(X)V(X)W_X, \quad (1)$$

where $f(x) \in [0, \infty)$ is the light contribution without shadows, and $V(x) \in [0, 1]$ is the shadow ray visibility. W_X is the contribution weight (i.e., estimate reciprocal PDF) given by

$$W_X = \frac{1}{\hat{p}_c(X)} \sum_i w_i, \quad (2)$$

where $\hat{p}_c(x) \propto f(x)$ is the target distribution at the current pixel c . In generalized resampled importance sampling [Lin et al. 2022],

SIGGRAPH Talks '24, July 27 - August 01, 2024, Denver, CO, USA

© 2024 Copyright held by the owner/author(s).

This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in *Special Interest Group on Computer Graphics and Interactive Techniques Conference Talks (SIGGRAPH Talks '24)*, July 27 - August 01, 2024, <https://doi.org/10.1145/3641233.3664320>.

the weight for a candidate sample is given by

$$w_i = m_i(T_i(x_i)) \hat{p}_c(T_i(x_i)) W_i \left| \frac{\partial T_i}{\partial x_i} \right|, \quad (3)$$

where $m_i(\cdot)$ is the weight of multiple importance sampling [Veach and Guibas 1995] that satisfies $\sum_i m_i(x) = 1$, $T_i(\cdot)$ is a bijective shift mapping from the candidate's domain to the integral domain Ω . Although ReSTIR reduces the variance in $f(X)W_X$, the visibility $V(X)$ can have a high variance (Fig. 1a). This variance is noticeable around shadow edges for a biased ReSTIR variant [Bitterli et al. 2020] that reuses the visibility of the initial sample for W_X . Our work reduces the shadow variance (Fig. 1b) by reusing spatiotemporal sample visibilities for $V(X)$.

3 OUR VISIBILITY REUSE

Our method reuses past visibilities while correcting the visibility of a spatial sample by tracing a shadow ray, similar to Wyman and Pantelev [2021]. Unlike Wyman and Pantelev, we estimate the visibility $\bar{V}_X \in [0, 1]$ of a selected sample X by the weighted average of spatiotemporal sample visibilities $\bar{V}_i \in [0, 1]$ as follows:

$$\bar{V}_X = \frac{\sum_i \bar{V}_i w_i}{\sum_i w_i}. \quad (4)$$

Since a sample X is selected according to w_i , this estimation is a variant of WIS. Although WIS is biased due to the normalization factor $\sum_i w_i$, this factor is eventually cancelled as follows:

$$\begin{aligned} \int_{\Omega} f(x)V(x)dx &\approx f(X)\bar{V}_X W_X = f(X) \frac{\sum_i \bar{V}_i w_i}{\sum_i w_i} \frac{\sum_i w_i}{\hat{p}_c(X)} \\ &= \frac{f(X)}{\hat{p}_c(X)} \sum_i \bar{V}_i w_i. \end{aligned} \quad (5)$$

The reused visibility \bar{V}_i is also estimated using the WIS variant (Eq. 4) in a chained manner, but its normalization factor is also recursively cancelled by $\bar{V}_i w_i$ as follows:

$$\bar{V}_i w_i = m_i(T_i(x_i)) \hat{p}_c(T_i(x_i)) \frac{\sum_k \bar{V}_k w_k}{\sum_k w_k} \frac{\sum_k w_k}{\hat{p}_c(x_i)} \left| \frac{\partial T_i}{\partial x_i} \right|, \quad (6)$$

where k is a pixel used to resample x_i . Therefore, if $\hat{p}_c(x) \propto f(x)$, our variance reduction technique is unbiased for static scenes (please see the supplementary document for details). If $\hat{p}_c(x) \not\propto f(x)$, the bias occurs depending on the approximation error of $\hat{p}_c(x)$. When $f(x)$ differs between RGB channels (e.g., colored light sources), $\hat{p}_c(x)$ can be approximated by the luminance of $f(x)$ for simplicity. In this case, the color of invisible lights can leak into shadows, though the luminance is unbiased. This color leak can be made less noticeable by combining Bitterli et al.'s visibility reuse [2020], because their method reduces the sampling of invisible lights. Since our technique is simple (Eq. 4), it is easy to implement on top of the existing visibility-reuse ReSTIR (please see the pseudocode in the supplementary document).

4 RESULTS

Fig. 1 shows the biased visibility-reuse ReSTIR [Bitterli et al. 2020; Wyman and Pantelev 2021] with our variance reduction technique. Our method significantly reduces undesirable shadow noise. Although our method can introduce a color leak bias around shadow

edges, this bias is barely noticeable even for colored light sources. For dynamic lighting, a shadow delay is inherited from the existing visibility reuse, but it is insignificant as shown in the supplementary video. Since our method improves the shadow quality with simple implementation, it is suitable for real-time rendering.

ACKNOWLEDGMENTS

We would like to thank the Amazon Lumberyard team for the BISTRO scene [2017].

REFERENCES

- Philippe Bekaert, Mateu Sbert, and Yves D. Willems. 2000. Weighted Importance Sampling Techniques for Monte Carlo Radiosity. In *EGWR '00*. 35–46.
- 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 (2020), 17 pages. <https://doi.org/10.1145/3386569.3392481>
- 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 (2022), 23 pages. <https://doi.org/10.1145/3528223.3530158>
- Amazon Lumberyard. 2017. Amazon Lumberyard Bistro, Open Research Content Archive (ORCA). <http://developer.nvidia.com/orca/amazon-lumberyard-bistro>
- Eric Veach and Leonidas J. Guibas. 1995. Optimally Combining Sampling Techniques for Monte Carlo Rendering. In *SIGGRAPH '95*. 419–428. <https://doi.org/10.1145/218380.218498>
- Chris Wyman and Alexey Pantelev. 2021. Rearchitecting Spatiotemporal Resampling for Production. In *HPG '21*. 19 pages. <https://doi.org/10.2312/hpg.20211281>