Efficient Visibility Reuse for Real-time ReSTIR (Supplementary Document)

YUSUKE TOKUYOSHI, Advanced Micro Devices, Inc., Japan

A SPATIOTEMPORAL RESERVOIR RESAMPLING (RESTIR)

Spatiotemporal reservoir resampling (ReSTIR) [Bitterli et al. 2020] is a chained form of the generalized *resampled importance sampling* (RIS) [Lin et al. 2022], and it increases the number of candidate samples by reusing spatiotemporal neighbor samples stored in each pixel. This algorithm first generates an initial candidate sample using classic RIS [Talbot 2005] for each pixel. Then, a sample X is resampled from spatiotemporally neighboring pixels i according to the weight w_i of each candidate sample. For the notations used in this document, please see Table 1. In generalized RIS, the weight for a candidate sample is given by

$$w_i = m_i(T_i(x_i))\hat{p}_c\left(T_i(x_i)\right) W_i \left|\frac{\partial T_i}{\partial x_i}\right|, \tag{1}$$

where $\hat{p}_c(x)$ is the target distribution approximately proportional to the integrand without the shadow ray visibility, $m_i(\cdot)$ is the weight of *multiple importance sampling* (MIS) [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 Ω_i to the integral domain Ω, and W_i is the contribution weight (estimated reciprocal PDF) for the candidate i. For a selected sample X, we get its contribution weight:

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

This contribution weight W_X will be used as W_i in the next resampling step. One high-quality MIS weight for RIS is Talbot MIS [2005], and it is generalized by Lin et al. [2022]. For ReSTIR, an accumulated candidate count M_i is used as a confidence weight for each distribution as follows:

$$m_i(x) = \frac{M_i \hat{p}_{\leftarrow i}(x)}{\sum_j M_j \hat{p}_{\leftarrow j}(x)}, \quad \text{where } \hat{p}_{\leftarrow i}(x) = \begin{cases} \hat{p}_i\left(T_i^{-1}(x)\right)\left|\partial T_i^{-1}/\partial x\right| & \text{if } T_i^{-1}(x) \text{ is defined} \\ 0 & \text{otherwise} \end{cases}. \tag{3}$$

We use the above MIS weight in our implementation. For other MIS weights, please refer to Lin et al. [2022].

B BIAS CANCELLATION FOR OUR VISIBILITY ESTIMATION

Our method estimates the visibility term based on a variant of *weighted importance sampling* (WIS) [Bekaert et al. 2000]:

$$\bar{V}_X = \frac{\sum_i \bar{V}_i w_i}{\sum_i w_i}. \tag{4}$$

WIS is a biased estimator due to the normalization factor $\sum_i w_i$. In our method, the reused visibility $\bar{V}_i$ in the numerator is also estimated using this WIS variant in a chained manner. Therefore, $\bar{V}_i$ has a bias unlike the original WIS method. However, the normalization factor for $\bar{V}_i$ is recursively cancelled by $\bar{V}_i w_i$ as follows:

$$\bar{V}_i w_i = \begin{cases} m_i(T_i(x_i))\hat{p}_c\left(T_i(x_i)\right) V(T_i(x_i)) W_i \left|\frac{\partial T_i}{\partial x_i}\right| & \text{if } x_i \text{ is an initial sample or spatial sample} \\ m_i(T_i(x_i))\hat{p}_c\left(T_i(x_i)\right) \frac{\sum_k \bar{V}_k w_k}{\cancel{\sum_k w_k}} \frac{\cancel{\sum_k w_k}}{\hat{p}_i(x_i)} \left|\frac{\partial T_i}{\partial x_i}\right| & \text{otherwise} \end{cases}, \tag{5}$$

Author's address: Yusuke Tokuyoshi, yusuke.tokuyoshi@amd.com, Advanced Micro Devices, Inc., Tokyo, Japan.

Table 1. Notations used in this document

Symbol	Description
c	Current pixel to compute the integral
i	Candidate index, or reused pixel
Ω	Domain of integration for the current pixel c
Ω_i	Domain of ith candidate sample
$f \quad : \Omega \rightarrow [0, \infty)$	Integrand without the shadow ray visibility, path contribution for lighting
$V \quad : \Omega \rightarrow [0, 1]$	Shadow ray visibility term in the integrand
$\hat{p}_c \quad : \Omega \rightarrow [0, \infty)$	Unnormalized target distribution approximately proportional to $f(x)$, $\mathrm{supp}(\hat{p}_c) \supseteq \mathrm{supp}(f) \cap \mathrm{supp}(V)$
$T_i \quad : \Omega_i \rightarrow \Omega$	Shift mapping from Ω_i to Ω
$X \quad \in \Omega$	Path sample selected via generalized RIS
$x_i \quad \in \Omega_i$	Path of candidate sample
$w_i \quad \in [0, \infty)$	Weight of candidate sample
$W_i \quad \in [0, \infty)$	Contribution weight, estimate reciprocal PDF
$M_i \quad \in [0, \infty)$	Confidence weight, accumulated count of candidate samples
$\bar{V}_i \quad \in [0, 1]$	Estimate visibility of ith candidate sample

where k is a pixel used to resample x_i. Therefore, if the contribution weight W_i is unbiased and lighting is static (i.e., $T_i(x_i) = x_i$ between frames), $\sum_i \bar{V}_i w_i$ is unbiased as follows:

$$\int_{\Omega} \hat{p}_c(x) V(x) \mathrm{d}x = \mathrm{E}\left[\sum_i \bar{V}_i w_i\right]. \tag{6}$$

For $\hat{p}_c(x) \propto f(x)$, this results in an unbiased visibility-reuse ReSTIR:

$$\int_{\Omega} f(x) V(x) \mathrm{d}x = \mathrm{E}\left[\frac{f(X)}{\hat{p}_c(X)} \sum_i \bar{V}_i w_i\right] = \mathrm{E}\left[f(X) \bar{V}_X W_X\right], \tag{7}$$

where $\frac{f(X)}{\hat{p}_c(X)}$ is constant.

C BIASED RESTIR WITH OUR VISIBILITY REUSE

Algorithm 1 shows the pseudocode for biased visibility-reuse ReSTIR [Bitterli et al. 2020; Wyman and Panteleev 2021] with our variance reduction technique. Bitterli et al.'s visibility reuse [2020] introduces a darkening bias in the contribution weight W_i, but it significantly reduces the variance. In addition, it mitigates a color leak bias for our method in the case where $\hat{p}_c(x) \not\propto f(x)$. The visibility $\bar{V}_i$ of a temporal neighbor is reused by assuming $T_i(x_i) = x_i$ between frames, while the visibility $\bar{V}_i$ of a spatial neighbor is corrected by tracing a shadow ray, as in Wyman and Panteleev [2021]. Omitting shadow rays for nearby spatial samples is an option that introduces a brightening bias (we do not use this biasing option in our experiments, i.e., $\mathsf{Distant}(c, i)$ is always true). To reduce the variance of the visibility $\bar{V}_X$ for a selected sample X, we estimate the visibility based on WIS (Eq. 4) instead of reusing the visibility of only one sample.

D EXPERIMENTAL RESULTS

Fig. 1 shows the visualization of the visibility term $\bar{V}_X$ for the previous method and our method. Our method reduces the variance of the visibility term while producing the same sample distribution as the previous method. On the other hand, for a scene with colored light sources, our method can introduce a small color leak bias around shadow edges (Fig. 2).

Algorithm 1: Biased ReSTIR with visibility reuse. Our variance reduction technique written in red is simple and easy to implement (only one line of code).

```
function ReSTIR(c)
    [x_c, W_c, M_c] ← RIS(c);
    V̄_c ← V(x_c);                                   // Trace a shadow ray for the initial sample
    if V̄_c = 0 then W_c ← 0;   // Visibility reuse for resampling weights [Bitterli et al. 2020]
    [x_c, W_c, M_c, V̄_c] ← TemporalResampling(c, [x_c, W_c, M_c, V̄_c]);
    [x_c, W_c, M_c, V̄_c] ← SpatialResampling(c, [x_c, W_c, M_c, V̄_c]);
    StoreReservoir(c, [x_c, W_c, M_c, V̄_c]);
    return f(x_c)V̄_c W_c;  // Reuse V̄_c instead of tracing a shadow ray [Wyman and Panteleev 2021]
end
function TemporalResampling(c, [x_c, W_c, M_c, V̄_c])
    i ← PickTemporalNeighbor(c);
    [x_i, W_i, M_i, V̄_i] ← LoadReservoir(i);
    M_i ← TemporalRejectionHeuristic(M_i, c, i);
    return Resample(c, [x_c, W_c, M_c, V̄_c], i, [x_i, W_i, M_i, V̄_i]);
end
function SpatialResampling(c, [x_c, W_c, M_c, V̄_c])
    i ← PickSpatialNeighbor(c);
    [x_i, W_i, M_i, V̄_i] ← LoadReservoir(i);
    M_i ← SpatialRejectionHeuristic(M_i, c, i);
    if Distant(c, i) then V̄_i ← V(T_i(x_i));    // Trace a shadow ray for a (distant) spatial sample
    return Resample(c, [x_c, W_c, M_c, V̄_c], i, [x_i, W_i, M_i, V̄_i]);
end
function Resample(c, [x_c, W_c, M_c, V̄_c], i, [x_i, W_i, M_i, V̄_i])
    [w_c, w_i] ← GetResamplingWeights(c, [x_c, W_c, M_c], i, [x_i, W_i, M_i]);           // Eq. 1
    w_sum ← w_c + w_i;
    ξ ← GenerateRandomNumber();
    if ξ < w_i/w_sum then X ← T_i(x_i) else X ← x_c ;
    W_X ← w_sum/p̂_c(X);                                                                  // Eq. 2
    M_X ← M_c + M_i;
    V̄_X ← (V̄_c w_c + V̄_i w_i)/w_sum;                            // Estimate the visibility (Eq. 4)
    return [X, W_X, M_X, V̄_X];
end
```

We can avoid this color leak bias by generating a sample for each color channel, although the computational cost is tripled.

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

(a) Previous method ($f(X)\bar{V}_X W_X$)

(b) Our method ($f(X)\bar{V}_X W_X$)

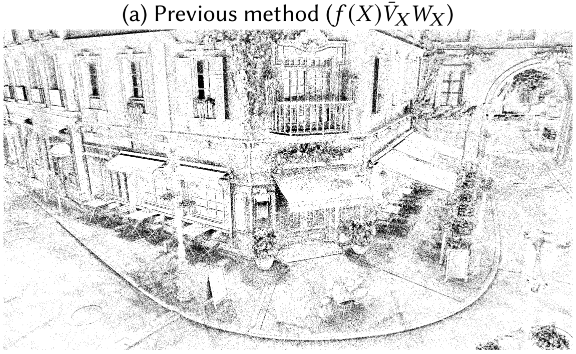

(c) Previous visibility term ($\bar{V}_X$)

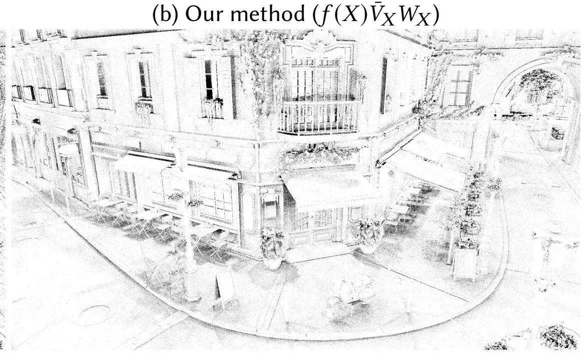

(d) Our visibility term ($\bar{V}_X$)

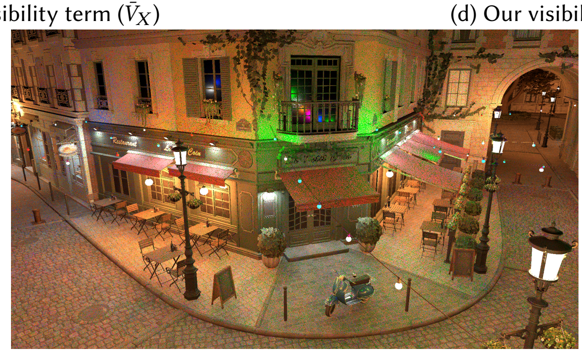

(e) No visibility term ($f(X)W_X$)

Fig. 1. Biased ReSTIR direct illumination (1920×1080 pixels, two rays per pixel, AMD Radeon™ RX 7900 XTX GPU) using the previous visibility reuse (a) and our visibility reuse (b). The images (c, d) visualize the visibility term $\bar{V}_X$. Our method reduces the undesirable noise in the visibility term.

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

Justin F. Talbot. 2005. *Importance Resampling for Global Illumination.* Master's thesis. Brigham Young U.

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 Panteleev. 2021. Rearchitecting Spatiotemporal Resampling for Production. In *HPG '21*. 19 pages. https://doi.org/10.2312/hpg.20211281

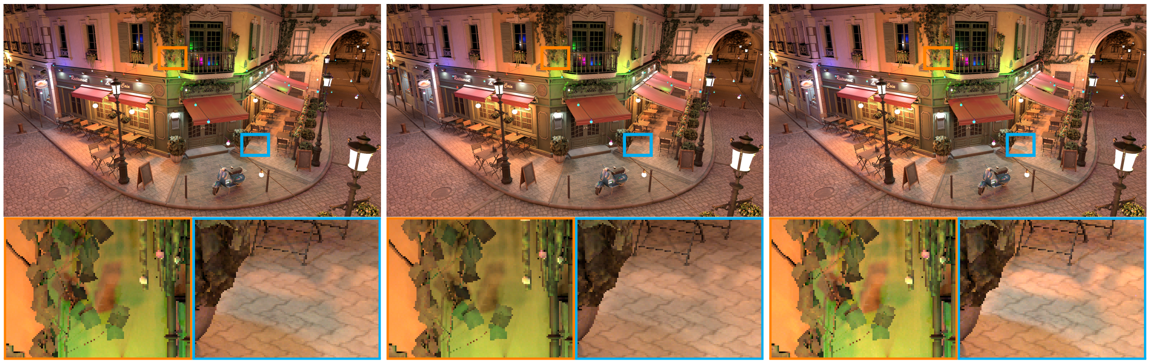

(a) Previous method (RMSPE: 15.3%) (b) Our method (RMSPE: 15.7%) (c) Reference

Fig. 2. Converged images using many samples for the biased ReSTIR with the previous visibility reuse (a) and our visibility reuse (b). Both images (a, b) have a darkening bias due to Bitterli et al.'s visibility reuse [2020]. For colored light sources, although our visibility estimation (b) further introduces a color leak bias around shadow edges, this color leak bias (0.4% in RMSPE) is significantly smaller than the darkening bias (15.3% in RMSPE).