Fast Volume Rendering with Spatiotemporal Reservoir Resampling (Supplemental Document)

DAQI LIN, University of Utah, USA CHRIS WYMAN, NVIDIA, USA CEM YUKSEL, University of Utah, USA

ACM Reference Format:

Daqi Lin, Chris Wyman, and Cem Yuksel. 2021. Fast Volume Rendering with Spatiotemporal Reservoir Resampling (Supplemental Document). *ACM Trans. Graph.* 40, 6, Article 279 (December 2021), 4 pages. https://doi.org/10.1145/3478513.3480499

1 RIS FOR PATH INTEGRALS

This subsection provides a more rigorous derivation of the RIS estimator of the path integral proposed in Section 3.3 in the paper. We assume the path tracing process being a random walk that generates *n* paths λ^{i} ($i \in \{1, ..., n\}$) that differ by lengths (or emission/scatter type in the volumetric case), responsible for *n* non-overlapping parts $(\int_{\Lambda^{i}} F(\lambda^{i}) d\lambda^{i})$ of the path integral $L = \int_{\Lambda} F(\lambda) d\lambda$ (where λ can be any length or type). In other words, $L = \sum_{i=1}^{\infty} \int_{\Lambda^i} F(\lambda^i) d\lambda^i$. An ordinary 1-sample Monte Carlo estimator simply sums the contributions from all sampled paths, i.e. $\langle L \rangle_{\text{MC}} = \sum_{i=1}^{n} \frac{F(\lambda^{i})}{p(\lambda^{i})}$. Note that *n* is a random variable and can pick any value from 0 to ∞ . We know that $\langle L \rangle_{\rm MC}$ is an unbiased estimator of the path integral *L*. However, we want to take advantage of RIS to avoid evaluating F for all paths. The traditional form of RIS assumes that all candidate samples are taken from the same sampling domain to estimate the same integral. Now, we want to use RIS to estimate a sum of integrals using candidate samples taken from the sampling domain of each integral, and only evaluate F for the chosen sample λ^r . This requires a slightly different definition of RIS (the "1/M" term is removed from the estimator). We now show that the RIS estimator

$$\langle L \rangle_{\rm ris} = \frac{F(\lambda^r)}{\hat{p}(\lambda^r)} \sum_{i=1}^n \frac{\hat{p}(\lambda^i)}{p(\lambda^i)}$$

where λ^r is chosen from $\lambda^1, ..., \lambda^n$ according to the weight $w(\lambda^i) = \frac{\hat{p}(\lambda^i)}{p(\lambda^i)}$ is also an unbiased estimator of the path integral, i.e. $\mathbb{E}[\langle L \rangle_{\text{ris}}] = \mathbb{E}[\langle L \rangle_{\text{MC}}]$. The proof goes as follows,

$$\mathbb{E}\left[\langle L \rangle_{\mathrm{ris}}\right] = \mathbb{E}\left[\mathbb{E}\left[\frac{F(\lambda^{r})}{\hat{p}(\lambda^{r})} | \lambda^{1}, ..., \lambda^{n}\right] \cdot \sum_{i=1}^{n} w(\lambda^{i})\right]$$

Authors' addresses: Daqi Lin, University of Utah, USA; Chris Wyman, NVIDIA, USA; Cem Yuksel, University of Utah, USA.

@ 2021 Copyright held by the owner/author (s). Publication rights licensed to ACM. 0730-0301/2021/12-ART279 \$15.00

https://doi.org/10.1145/3478513.3480499

$$= \mathbb{E}\left[\sum_{j=1}^{n} \left(\frac{F(\lambda^{j})}{\hat{p}(\lambda^{j})} \cdot \frac{w(\lambda^{j})}{\sum_{i=1}^{n} w(\lambda^{i})}\right) \cdot \sum_{i=1}^{n} w(\lambda^{i})\right]$$
$$= \mathbb{E}\left[\sum_{j=1}^{n} \frac{F(\lambda^{j})}{\hat{p}(\lambda^{j})} \cdot \frac{\hat{p}(\lambda^{j})}{p(\lambda^{j})}\right]$$
$$= \mathbb{E}\left[\sum_{j=1}^{n} \frac{F(\lambda^{j})}{p(\lambda^{j})}\right] = \mathbb{E}[\langle I \rangle_{MC}]. \quad \Box$$

2 COMPARISON WITH VERTEX REUSE

Section 4.2 in the paper notes we can either resample paths by reusing path vertices \mathbf{x}_i or by reusing directions $\boldsymbol{\omega}_i$. We chose to reuse directions; while costs are somewhat higher for direction reuse, it reduces noise fairly significantly. This is because paths act a little like virtual point lights under reuse, which leads to more singularities and fireflies. However, if these could be reduced, it may make sense to switch back to vertex reuse for the improved performance. See Figure 1 to compare visually between vertex and directional reuse.

3 DENOISED RESULT

In Figure 2, we compare denoising applied to both our baseline and our resampling technique; both using the new OptiX 7.3 temporal denoising mode. While OptiX produces amazingly denoised results in both cases, the better sampling provided by our technique preserves much higher frequency details in the animation, while the baseline gives a smoother, more washed out look. Part of this is also due to the improved motion vectors we provide with our novel temporal reprojection plus velocity resampling.

Please see the supplementary video to compare the denoised results under animation.

4 PSEUDOCODE

We provide the pseudocode of our fast volume rendering with spatiotemporal reservoir resampling (Volumetric ReSTIR) in this supplemental document. Algorithm 1 gives an outline of the full algorithm, ignoring surface bounces for simplicity. Algorithm 2 shows the reservoir class we used for storing and streaming the samples. Algorithm 3, Algorithm 5, and Algorithm 6 shows four individual stages of our algorithm: initial sampling, temporal reuse, spatial reuse, and final shading. We provide the detail of how we resample the target function \hat{p} in Algorithm 4.

ACM Trans. Graph., Vol. 40, No. 6, Article 279. Publication date: December 2021.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

279:2 • Lin, Wyman, and Yuksel



Fig. 1. Compare vertex and direction reuse. (We reuse directions.) Here we show 7-bounce multiple scattering in the Bunny Cloud scene. While vertex reuse provides a slight performance win, it introduces additional sampling variance due to large weights of potentially irrelevant paths during spatial reuse.



Baseline (1 bounce) with denoising

Ours (1 bounce) with denoising



Baseline (7 bounces) with denoising

Ours (7 bounces) with denoising

Fig. 2. The OptiX 7.3 temporal denoiser applied to both our baseline and our method, using (top) single and (bottom) 7-bounce multiple scattering.

ACM Trans. Graph., Vol. 40, No. 6, Article 279. Publication date: December 2021.

Algorithm 1: Full algorithm.

-		
<pre>1 function VolumetricReSTIR()</pre>		
2	allocate an image size array of reservoirs	
3	foreach $q \in Image$ do	
4	$reservoirs[q] \leftarrow InitialSampling(q)$	
5	end	
6	foreach $q \in Image$ do	
7	$reservoirs[q] \leftarrow TemporalReuse(q)$	
8	end	
9	foreach $q \in Image$ do	
10	$reservoirs[q] \leftarrow SpatialReuse(q)$	
11	end	
12	prevFrameReservoirs ← reservoirs	
13	foreach $q \in Image$ do	
14	pixel color \leftarrow FinalShading(q)	
15	end	

Algorithm 2: Pseudocode of the Reservoir class.
1 class Reservoir
$_{2}$ $\lambda \leftarrow \text{null}$
$/\!/$ the content of the path sample can be seen in line 3, Algorithm 4.
$w_{sum} \leftarrow 0$
$4 \qquad M \leftarrow 0$
$\hat{p} \leftarrow 0$
6 function update $(p_i, \hat{p}_i, \lambda_i)$
7 if null sample then
$8 \qquad M \leftarrow M + 1$
9 else
10 $w_i \leftarrow \hat{p}_i/p_i$
11 $w_{sum} \leftarrow w_{sum} + w_i$
12 $M \leftarrow M + 1$
// Generate a random number ξ .
13 if $\xi < w_i/w_{sum}$ then
14 $\lambda \leftarrow \lambda_i$
15 $\hat{p} \leftarrow \hat{p}_i$
function combineReservoir(\hat{p}_i, r_i, m)
17 $w_i \leftarrow \hat{p}_i / (r_i \cdot \hat{p}) \cdot r_i \cdot w_{sum} \cdot m$
18 $w_{sum} \leftarrow w_{sum} + w_i$
$M \leftarrow M + r_i.M$
20 if $\xi < w_i / w_{sum}$ then
21 $\lambda \leftarrow \lambda_i$
22 $\hat{p} \leftarrow \hat{p}_i$

Algorithm 3: Pseudocode of the initial sampling pass.		
1 function InitialSampling(q)		
2 Reservoir r		
$[\mathbf{x}_0, \boldsymbol{\omega}_0] \leftarrow$ ray origin and direction from pixel q		
// <i>M</i> is the number of random walks		
4 for $m \leftarrow 1$ to M do		
5 Reservoir r _m		
$6 \qquad \qquad \hat{p} \leftarrow 1$		
7 $p \leftarrow 1$		
8 $\lambda_0 \leftarrow$ Path with only camera vertex		
9 for $i \leftarrow 1$ to K do		
10 $[z_i, \sigma_t^*(\mathbf{x}_i), T^*] \leftarrow \text{RegularTracking}(\mathbf{x}_{i-1}, \omega_{i-1})$		
11 $\lambda_i \leftarrow \text{AddPathVertex}(\lambda_{i-1}, \mathbf{x}_i)$		
12 if $z_i == z_s$ then		
13 if $i == 1$ then		
$14 \qquad p \leftarrow p \cdot T^*$		
15 $\hat{p} \leftarrow \hat{p} \cdot L^s T^*$		
16 $r_m.update(p, \hat{p}, \lambda_i)$		
17 else		
18 $r_m.update(null sample)$		
19 break		
20 if $\sigma_t(\mathbf{x}_i) == 0$ then		
21 $r_m.update(null sample)$		
22 break		
$p \leftarrow p \cdot \sigma_t^*(\mathbf{x}_i) T^*$		
$24 \qquad \hat{p} \leftarrow \hat{p} \cdot T^*$		
25 if volume contains emission then		
26 $\hat{p'} \leftarrow \hat{p} \cdot \sigma_a(\mathbf{x}_i) L_e^m(\mathbf{x}_i)$		
// the superscript e marks an emission path		
27 $r_m.update(p, p', \lambda_i^e)$		
// Scattering event. (p_{NEE} is the pdf of the light sample)		
28 $[L_s, \mathbf{x}_{i+1}, \boldsymbol{\omega}_i, p_{\text{NEE}}] \leftarrow \text{SampleLight}(\mathbf{x}_i)$		
// ray marching with a larger step size \tilde{T} (Estimate Values at the Shadow ($\pi - \pi'$)		
$\begin{array}{c} 1 \leftarrow \text{Estimatevolumetriconadow}(\mathbf{x}_i, \mathbf{x}_{i+1}) \\ \hat{\mathbf{x}}_i \leftarrow \hat{\mathbf{x}}_i = \mathbf{x}_i (\mathbf{x}_i) \mathbf{x}_i (\mathbf{x}_i - \mathbf{x}_{i+1}) \\ \hat{\mathbf{x}}_i \leftarrow \hat{\mathbf{x}}_i = \mathbf{x}_i (\mathbf{x}_i) \mathbf{x}_i (\mathbf{x}_i - \mathbf{x}_{i+1}) \\ \hat{\mathbf{x}}_i \leftarrow \hat{\mathbf{x}}_i = \mathbf{x}_i (\mathbf{x}_i) \mathbf{x}_i (\mathbf{x}_i - \mathbf{x}_i) \\ \hat{\mathbf{x}}_i = \mathbf{x}_i (\mathbf{x}_i) \mathbf{x}_i (\mathbf{x}_i - \mathbf{x}_i) \\ \hat{\mathbf{x}}_i = \mathbf{x}_i (\mathbf{x}_i) \mathbf{x}_i (\mathbf{x}_i - \mathbf{x}_i) \\ \hat{\mathbf{x}}_i = \mathbf{x}_i (\mathbf{x}_i) \mathbf{x}_i (\mathbf{x}_i - \mathbf{x}_i) \\ \hat{\mathbf{x}}_i = \mathbf{x}_i (\mathbf{x}_i) \mathbf{x}_i (\mathbf{x}_i - \mathbf{x}_i) \\ \hat{\mathbf{x}}_i = \mathbf{x}_i (\mathbf{x}_i) \mathbf{x}_i (\mathbf{x}_i - \mathbf{x}_i) \\ \hat{\mathbf{x}}_i = \mathbf{x}_i (\mathbf{x}_i) \mathbf{x}_i (\mathbf{x}_i) \mathbf{x}_i (\mathbf{x}_i) \\ \hat{\mathbf{x}}_i = \mathbf{x}_i (\mathbf{x}_i) \mathbf{x}_i (\mathbf{x}_i) \mathbf{x}_i (\mathbf{x}_i) \\ \hat{\mathbf{x}}_i = \mathbf{x}_i (\mathbf{x}_i) \mathbf{x}_i (\mathbf{x}_i) \mathbf{x}_i (\mathbf{x}_i) \\ \hat{\mathbf{x}}_i = \mathbf{x}_i (\mathbf{x}_i) \mathbf{x}_i (\mathbf{x}_i) \mathbf{x}_i (\mathbf{x}_i) \mathbf{x}_i (\mathbf{x}_i) \\ \hat{\mathbf{x}}_i = \mathbf{x}_i (\mathbf{x}_i) \mathbf{x}_i (\mathbf{x}_i) \mathbf{x}_i (\mathbf{x}_i) \mathbf{x}_i (\mathbf{x}_i) \mathbf{x}_i (\mathbf{x}_i) \\ \hat{\mathbf{x}}_i = \mathbf{x}_i (\mathbf{x}_i) \mathbf{x}_i (x$		
$\begin{array}{c c} \mathbf{y} \leftarrow p \cdot \sigma_{s}(\mathbf{x}_{i}) \rho(\mathbf{x}_{i}, -\boldsymbol{\omega}_{k-1}, \boldsymbol{\omega}_{i}) L_{s} \mathbf{I} \\ \mathbf{y} \leftarrow p \cdot \sigma_{s}(\mathbf{x}_{i}) \rho(\mathbf{x}_{i}, -\boldsymbol{\omega}_{k-1}, \boldsymbol{\omega}_{i}) L_{s} \mathbf{I} \end{array}$		
$\begin{array}{c c} y \leftarrow p \cdot p_{\text{NEE}} \\ y \leftarrow p \cdot p \cdot p \cdot p_{\text{NEE}} \\ y \leftarrow p \cdot p \cdot p \cdot p \cdot p_{\text{NEE}} \\ y \leftarrow p \cdot p$		
$r_m.update(p, p, AddPatnvertex(\lambda_i, \mathbf{x}_{i+1}))$		
// Prepare for the next bounce.		
$ \begin{array}{c} \omega_{l} \leftarrow \text{sampler naser unction}(\mathbf{x}_{l}) \\ \hat{\mathbf{b}} \leftarrow \hat{\mathbf{b}} \cdot \boldsymbol{\sigma} (\mathbf{x}_{l}) \hat{\mathbf{c}}(\mathbf{x}_{l} \leftarrow \boldsymbol{\omega}_{l}) \end{array} $		
$\begin{array}{c c} y \leftarrow p \cdot o_{\mathcal{S}}(\mathbf{x}_i) \rho(\mathbf{x}_i, -\boldsymbol{\omega}_{i-1}, \boldsymbol{\omega}_i) \\ p \leftarrow p \cdot o_{\mathcal{S}}(\mathbf{x}_i) \rho(\mathbf{x}_i, -\boldsymbol{\omega}_{i-1}, \boldsymbol{\omega}_i) \end{array}$		
$\sum_{i=1}^{35} p \leftarrow p \cdot p(\mathbf{x}_i, -\boldsymbol{\omega}_{i-1}, \boldsymbol{\omega}_i)$		
30 CIIU		
r_{m} $M \leftarrow 1$		
$\frac{1}{28} = \frac{1}{r} \operatorname{combineReservoir}(\hat{\rho}(r \lambda) r 1)$		
and $p(r_m, \lambda), r_m, 1$		
40 return r		
iv icculii /		

ACM Trans. Graph., Vol. 40, No. 6, Article 279. Publication date: December 2021.

Algorithm 4: Resample target function.

1 f	1 function ResampleTargetFunction(λ , q)		
2	$[\mathbf{x}_0, \boldsymbol{\omega}_0] \leftarrow$ ray origin and direction from pixel q		
	// k' is the number of scattering events (plus 1 for emissive paths ^1). $\mathbf{x}_{k'+1}$		
	and $\omega_{k'}$ are the position and direction of the light sample.		
3	$[k', is Emission, z_1, \omega_1, z_2, \omega_2,, z_{k'}, \omega_{k'}, \mathbf{x}_{k'+1}] \leftarrow \lambda$		
4	$\hat{p} \leftarrow 1$		
5	for $i \leftarrow 1$ to k' do		
6	$\mathbf{x}_i = \mathbf{x}_{i-1} + z_i \boldsymbol{\omega}_{i-1}$ // Direction reuse		
7	$\tilde{T} \leftarrow \text{RayMarching}(\mathbf{x}_{i-1}, \mathbf{x}_i)$		
8	if $z_i == z_s$ then		
9	if $i == 1$ then		
10	$\hat{p} \leftarrow L^s \tilde{T}$		
11	return \hat{p}		
12	else		
13	return 0		
14	if isEmission then		
15	return $\hat{p} \cdot \sigma_a(\mathbf{x}_i) \tilde{T} L_e^m(\mathbf{x}_i)$		
16	$\hat{p} \leftarrow \hat{p} \cdot \sigma_s(\mathbf{x}_i)\tilde{T}$		
17	$\hat{p} \leftarrow \hat{p} \cdot \rho(\mathbf{x}_i, -\boldsymbol{\omega}_{i-1}, \boldsymbol{\omega}_i)$		
18	end		
	// Transmittance of the NEE segment		
19	$\tilde{T} \leftarrow \text{RayMarching}(\mathbf{x}_{k'}, \mathbf{x}_{k'+1})$		
20	$\hat{p} \leftarrow \hat{p} \cdot \tilde{T}$		
21	return \hat{p}		

Alg	orithm 5: Pseudocode of spatial/temporal reuse.
1 fu	nction TemporalReuse(q)
2	$q' \leftarrow \text{TemporalReprojection}(q, \text{reservoirs}[q])$
3	return CombineReservoirs(reservoirs[q],
	prevFrameReservoirs[q'], q, q'})
4 fu	<pre>nction SpatialReuse(q)</pre>
5	$S \leftarrow \text{pickSpatialNeighbors}(q)$
6	return CombineReservoirs(reservoirs[q], {reservoir[q']
	$q' \in S \}, q, \{q' \in S\})$
7 fu	nction CombineReservoirs $(r_0, r_1,, r_N, q_0, q_1,, q_N)$
8	Reservoir s
9	foreach $r_i \in \{r_0,, r_N\}$ do
10	$p_{\text{sum}} \leftarrow 0$
11	$k \leftarrow 0$
	// Compute MIS weight
12	foreach $q_s \in \{q_0,, q_N\}$ do
13	$\hat{p}_{q_s}(r_i.\boldsymbol{\lambda}) \leftarrow \text{ResampleTargetFunction}(r_i.\boldsymbol{\lambda}, q_s)$
14	$p_{\text{sum}} \leftarrow p_{\text{sum}} + \hat{p}_{q_s}(r_i.\boldsymbol{\lambda}) \cdot r_i.M$
15	$k \leftarrow k + r_i.M$
16	end
17	$m \leftarrow rac{\hat{p}_{q_i}(r_i.\lambda)}{p_{\mathrm{sum}}/k}$ // MIS weight
18	$\hat{p}_q(r_i.\boldsymbol{\lambda}) \leftarrow \text{ResampleTargetFunction}(r_i.\boldsymbol{\lambda}, q_0)$
19	s.combineReservoir($\hat{p}_q(r_i.\lambda), r_i, m$)
20	end
21	return s

Algorithm 6: Final shading.		
<pre>1 function FinalShading(q)</pre>		
2	$r \leftarrow \text{reservoirs}[q]$	
	// Analytical transmittance values are computed.	
3	$f(r.\lambda) \leftarrow \text{compute integrand}$	
4	return $\frac{f(r.\lambda)}{\hat{p}(r.\lambda)} \frac{r.w_{sum}}{r.M}$	

¹Directly visible surface emission also counts as an emission path, with the corresponding terms ("if emission path") in Equation 15 and 25 in the paper changed to $L^{s}(\mathbf{x}_{1} \rightarrow \mathbf{x}_{0})$.

ACM Trans. Graph., Vol. 40, No. 6, Article 279. Publication date: December 2021.