

# ReSTIR PT Enhanced: Algorithmic Advances for Faster and More Robust ReSTIR Path Tracing – Supplemental Document

DAQI LIN, NVIDIA, USA

MARKUS KETTUNEN, NVIDIA, Finland

CHRIS WYMAN, NVIDIA, USA

CCS Concepts: • **Computing methodologies** → **Ray tracing**.

Additional Key Words and Phrases: Optimization, Sampling, Light Transport, Real Time

## ACM Reference Format:

Daqi Lin, Markus Kettunen, and Chris Wyman. 2026. ReSTIR PT Enhanced: Algorithmic Advances for Faster and More Robust ReSTIR Path Tracing – Supplemental Document. *Proc. ACM Comput. Graph. Interact. Tech.* 9, 1 (May 2026), 13 pages. <https://doi.org/10.1145/3804494>

## 1 LOBE-INDEXED PATHS AND PRIMARY SAMPLE SPACE

Modern path tracers often evaluate only a single sampled lobe for each path vertex, effectively splitting the integrand associated with the same geometric path across all lobe combinations. In ReSTIR PT, this separation allows more effective path reuse, as instead of classifying the *whole vertex* as either rough or (near-)specular, leading to suboptimal reuse decisions, only (near-)specular *lobes* are postponed with random replay.

In addition to sampling a single BSDF lobe to extend the path, each vertex also performs next event estimation (NEE) with all lobes enabled. Therefore, the path space integral can be written as

$$I = \sum_{d=1}^{\infty} \sum_{\bar{\ell} \in L_d} \int_{\Omega_d} \omega_{n(\bar{\ell})}(\bar{\mathbf{x}}) f_{\bar{\ell}}(\bar{\mathbf{x}}) d\bar{\mathbf{x}}, \quad (1)$$

where each path  $\bar{\mathbf{x}}$  now contains a lobe sequence  $\bar{\ell} = (\ell_j)_{j=1}^{d-1}$  with  $\ell_j \in \{0, 1, 2, \dots, N_{\text{lobe}}\}$  defining the lobe used at vertex  $\mathbf{x}_j$ , with a NEE sample (at the last vertex  $\mathbf{x}_d$ ) encoded by index 0.  $N_{\text{lobe}}$  is the maximum number of BSDF lobes in a vertex,  $L_d$  is the set of all possible lobe sequences, and  $n(\bar{\ell}) \in \{1, 2\}$  differentiates between BSDF and NEE sampling (based on  $\ell_{d-1}$ ) for MIS weights like  $\omega_t(\bar{\mathbf{x}}) = \frac{p_t(\bar{\mathbf{x}})}{p_1(\bar{\mathbf{x}}) + p_2(\bar{\mathbf{x}})}$ . Written out explicitly, the lobe-specific path contribution is

$$f_{\bar{\ell}}(\bar{\mathbf{x}}) = \prod_{j=0}^{d-1} \left( \rho_{\ell_j}(\mathbf{x}_{j+1} \rightarrow \mathbf{x}_j \rightarrow \mathbf{x}_{j-1}) G(\mathbf{x}_j \leftrightarrow \mathbf{x}_{j+1}) V(\mathbf{x}_j \leftrightarrow \mathbf{x}_{j+1}) \right) L_e(\mathbf{x}_d \rightarrow \mathbf{x}_{d-1}), \quad (2)$$

where  $\rho_{\ell_j}$  is the chosen BSDF lobe  $\ell_j$  at  $\mathbf{x}_j$  for vertices  $j \geq 1$ , and the sensor response  $W_e(\mathbf{x}_1 \rightarrow \mathbf{x}_0)$  for the sensor vertex  $j = 0$ .

In principle, this allows a unique mapping to primary sample spaces (PSS) can then be established, turning the  $(\bar{\mathbf{x}}, \bar{\ell})$  pair into a random number sequence  $\bar{\mathbf{u}} \in \mathcal{U}_{t,d}$  where  $d$  (path length) and  $t \in \{1, 2\}$

---

Authors' addresses: Daqi Lin, [daqil@nvidia.com](mailto:daqil@nvidia.com), NVIDIA, USA; Markus Kettunen, [mkettunen@nvidia.com](mailto:mkettunen@nvidia.com), NVIDIA, Finland; Chris Wyman, [chris.wyman@acm.org](mailto:chris.wyman@acm.org), NVIDIA, USA.

---

Please use nonacm option or ACM Engage class to enable CC licenses



This work is licensed under a Creative Commons Attribution 4.0 International License.

© 2026 Copyright held by the owner/author(s).

ACM 2577-6193/2026/5-ART

<https://doi.org/10.1145/3804494>

(BSDF or NEE light sampling) specifies the dimensionality of the primary sample space, giving the integral

$$I = \sum_{d=1}^{\infty} \sum_{t=1}^2 \int_{\mathcal{U}_{t,d}} F(\bar{\mathbf{u}}) d\bar{\mathbf{u}} \quad (3)$$

where  $F(\bar{\mathbf{u}}) = \frac{\omega_{n(\bar{t})}(\mathcal{X}_{t,d}(\bar{\mathbf{u}}))f_{\bar{t}}(\mathcal{X}_{t,d}(\bar{\mathbf{u}}))}{p_t(\mathcal{X}_{t,d}(\bar{\mathbf{u}}))}$  is the PSS integrand and  $\mathcal{X}_{t,d}$  is a function that warps the random sequence into the lobe-indexed path. Note that  $p_t(\mathcal{X}_{t,d}(\bar{\mathbf{u}}))$  is the path-space sampling PDF in the path tracer that accounts for the change of variables between the path space and the PSS. The PSS source PDF  $p(\bar{\mathbf{u}}) = 1$  for paths directly produced by the path tracer.

With PSS parameterization, the Jacobian determinant for the hybrid shift can be expressed as

$$\left| \frac{\partial \bar{\mathbf{u}}^y}{\partial \bar{\mathbf{u}}^x} \right| = \frac{p_{k-1}^y(\omega'_{k-1}, \ell_{k-1})G(\mathbf{y}_{k-1} \rightarrow \mathbf{x}_k)p_k^y(\omega_k, \ell_k)}{p_{k-1}^x(\omega_{k-1}, \ell_{k-1})G(\mathbf{x}_{k-1} \rightarrow \mathbf{x}_k)p_k^x(\omega_k, \ell_k)}, \quad (4)$$

where  $p_k^x(\omega_k, \ell_k) \equiv p(\omega_k, \ell_k | \mathbf{x}_k, -\omega_{k-1})$  is the joint PDF of sampling the direction  $\omega_k = \widehat{\mathbf{x}_k \mathbf{x}_{k+1}} = \frac{\mathbf{x}_k \mathbf{x}_{k+1}}{\|\mathbf{x}_k \mathbf{x}_{k+1}\|}$  and lobe  $\ell_k$  at  $\mathbf{x}_k$  with outgoing direction  $-\omega_{k-1}$ . Other terms like  $p_{k-1}^x(\omega_{k-1}, \ell_{k-1})$  are similarly defined.  $\omega'_{k-1}$  is a shorthand for  $\widehat{\mathbf{y}_{k-1} \mathbf{x}_k}$ , the direction formed by reconnection. The single-sided geometry term  $G(\mathbf{x} \rightarrow \mathbf{y})$  is defined as  $\cos \theta / \|\mathbf{x} - \mathbf{y}\|^2$  where  $\theta$  is the angle made by  $\widehat{\mathbf{y}_{k-1} \mathbf{x}_k}$  and the normal at  $\mathbf{y}_{k-1}$ . Note that the Jacobian is simply the result of reconnection. Random replay is an identity shift in PSS, therefore having a Jacobian of 1.

## 2 DETAILED DERIVATION OF THE FOOTPRINT THRESHOLDS

The dual footprint threshold is

$$(\max(p_{k-1}^x(\omega_{k-1})G(\mathbf{x}_{k-1} \rightarrow \mathbf{x}_k), p_k^x(\omega_k)G(\mathbf{x}_k \rightarrow \mathbf{x}_{k-1})))^{-1} \geq \frac{c}{100} \frac{\|\mathbf{x}_0 - \mathbf{x}_1\|^2}{\langle n_{\mathbf{x}_1}, \widehat{\mathbf{x}_1 \mathbf{x}_0} \rangle / (4\pi)}, \quad (5)$$

where  $p_{k-1}^x(\omega_{k-1})$  and  $p_k^x(\omega_k)$  are PDFs of sampling  $\omega_{k-1}$  and  $\omega_k$  at  $\mathbf{x}_{k-1}$  and  $\mathbf{x}_k$ , respectively. The analysis assumes samples without lobe indices, with marginal PDF for direction sampling available. The lobe-indexed case is discussed in Section 3.

We consider shifting a path  $\bar{\mathbf{x}}$  starting from pixel A to a path  $\bar{\mathbf{y}}$  in a neighboring pixel B. For any pair of corresponding primary hits  $\mathbf{x}_1$  and  $\mathbf{y}_1$ , we assume their world-space distance is bounded by the primary ray footprint radius  $R_{\text{pri}}^{\bar{\mathbf{x}}}$  (“spread at the primary vertex” [Müller et al. 2021]) up to a constant factor,

$$|\mathbf{x}_1 - \mathbf{y}_1| < C_1 R_{\text{pri}}^{\bar{\mathbf{x}}}, \quad R_{\text{pri}}^{\bar{\mathbf{x}}} = \sqrt{\frac{\|\mathbf{x}_0 - \mathbf{x}_1\|^2}{\langle n_{\mathbf{x}_1}, \widehat{\mathbf{x}_1 \mathbf{x}_0} \rangle / (4\pi)}}, \quad (6)$$

and similarly  $|\mathbf{x}_1 - \mathbf{y}_1| \leq C_1 R_{\text{pri}}^{\bar{\mathbf{y}}}$ . Before reconnection, we further assume that random replay produces corresponding vertices  $\mathbf{x}_i$  and  $\mathbf{y}_i$  on the same (locally flat) surface and that the distances between them grow at most linearly:

$$|\mathbf{x}_i - \mathbf{y}_i| < C_2 |\mathbf{x}_1 - \mathbf{y}_1|. \quad (7)$$

These bounds reflect that neighboring pixels cannot diverge arbitrarily in a finite scene.

For a potential reconnection vertex  $\mathbf{x}_k$  on  $\bar{\mathbf{x}}$ , connecting from  $\mathbf{y}_{k-1}$  instead of  $\mathbf{x}_{k-1}$  perturbs the Jacobian of the shift  $T$ . A sufficient condition for a stable mapping is that the relative change in the area density at  $\mathbf{x}_k$  and the relative change in the BSDF sampling density at the next vertex are both

small. We express this as

$$\left\{ \begin{array}{l} \left| \frac{p_{k-1}^y(\omega'_{k-1})G(\mathbf{y}_{k-1} \rightarrow \mathbf{x}_k) - p_{k-1}^x(\omega_{k-1})G(\mathbf{x}_{k-1} \rightarrow \mathbf{x}_k)}{p_{k-1}^x(\omega_{k-1})G(\mathbf{x}_{k-1} \rightarrow \mathbf{x}_k)} \right| < \epsilon, \\ \left| \frac{p_k^y(\omega_k) - p_k^x(\omega_k)}{p_k^x(\omega_k)} \right| < \epsilon, \end{array} \right. \quad (8)$$

for a small  $\epsilon$ , which ensures  $(1 - \epsilon)^2 < |\partial T / \partial \bar{\mathbf{u}}^x| < (1 + \epsilon)^2$ . Our goal is to turn these conditions into simple, local thresholds based on ray footprints and inverse ray footprints.

*Ray footprint threshold.* The term  $p_{k-1}^x(\omega_{k-1})G(\mathbf{x}_{k-1} \leftrightarrow \mathbf{x}_k)$  is the area density of the vertex  $\mathbf{x}_k$  when the path  $\bar{\mathbf{x}}$  is generated by path tracing. We denote it by

$$p_k^{\bar{\mathbf{x}}}(\mathbf{x}_k) \equiv p_{k-1}^x(\omega_{k-1})G(\mathbf{x}_{k-1} \leftrightarrow \mathbf{x}_k), \quad (10)$$

and refer to its inverse  $1/p_k^{\bar{\mathbf{x}}}(\mathbf{x}_k)$  as the *ray footprint* at  $\mathbf{x}_k$ . Intuitively, this footprint measures the typical spacing between nearby area samples. Analogously, tracing from  $\mathbf{y}_{k-1}$  induces

$$p_k^{\bar{\mathbf{y}}}(\mathbf{x}_k) \equiv p_{k-1}^y(\omega'_{k-1})G(\mathbf{y}_{k-1} \leftrightarrow \mathbf{x}_k), \quad (11)$$

and the left-hand side of (8) becomes

$$\frac{|p_k^{\bar{\mathbf{y}}}(\mathbf{x}_k) - p_k^{\bar{\mathbf{x}}}(\mathbf{x}_k)|}{p_k^{\bar{\mathbf{x}}}(\mathbf{x}_k)}. \quad (12)$$

Directly comparing  $p_k^{\bar{\mathbf{x}}}$  and  $p_k^{\bar{\mathbf{y}}}$  is awkward. Instead, we make the empirical assumption that random replay preserves the area density along  $\bar{\mathbf{y}}$ , i.e.,

$$p_k^{\bar{\mathbf{x}}}(\mathbf{x}_k) \approx p_k^{\bar{\mathbf{y}}}(T_{\text{replay}}(\mathbf{x}_k)), \quad (13)$$

where  $T_{\text{replay}}$  denotes continuing random replay from  $\mathbf{y}_{k-1}$  to produce a vertex near  $\mathbf{x}_k$ . Under this approximation, (8) is controlled by the variation of a single density function  $p_k^{\bar{\mathbf{y}}}(x)$  under a small spatial shift:

$$\frac{|p_k^{\bar{\mathbf{y}}}(\mathbf{x}_k) - p_k^{\bar{\mathbf{y}}}(T_{\text{replay}}(\mathbf{x}_k))|}{p_k^{\bar{\mathbf{y}}}(T_{\text{replay}}(\mathbf{x}_k))}. \quad (14)$$

We now assume a footprint-based smoothness condition on the area density: for any area sample  $x$  with density  $p_{\text{area}}(x)$ , the relative change of  $p_{\text{area}}$  inside a disk of radius proportional to the footprint  $\sqrt{1/p_{\text{area}}(x)}$  is bounded by  $\epsilon$ . Formally, there exists a constant  $c_2$  such that if  $y, z$  satisfy

$$\max(\|y - x\|, \|z - x\|) < \sqrt{\frac{c_2}{p_{\text{area}}(x)}}, \quad (15)$$

then

$$\frac{|p_{\text{area}}(z) - p_{\text{area}}(y)|}{p_{\text{area}}(y)} < \epsilon. \quad (16)$$

In words, the density is approximately constant within a footprint-sized neighborhood.

Since (according to Equation 7) the distance between  $\mathbf{x}_k$  and  $T_{\text{replay}}(\mathbf{x}_k)$  can be bounded by the primary footprint,

$$|T_{\text{replay}}(\mathbf{x}_k) - \mathbf{x}_k| \leq c_1 R_{\text{pri}}^{\bar{\mathbf{y}}} \quad \text{where } c_1 = C_1 C_2. \quad (17)$$

Combining this geometric bound, the smoothness assumption (Equation 15, 16), and all the derivations above, a sufficient condition for (8) is

$$c_1 R_{\text{pri}}^{\bar{y}} < \sqrt{\frac{c_2}{p_k^{\bar{y}}(\mathbf{x}_k)}} \implies \frac{1}{p_k^{\bar{y}}(\mathbf{x}_k)} > \frac{c_1^2}{c_2} (R_{\text{pri}}^{\bar{y}})^2. \quad (18)$$

By shift invertibility, the same condition must hold for  $\bar{\mathbf{x}}$ , yielding

$$(p_{k-1}^{\bar{x}}(\omega_{k-1})G(\mathbf{x}_{k-1} \rightarrow \mathbf{x}_k))^{-1} = \frac{1}{p_k^{\bar{x}}(\mathbf{x}_k)} > \frac{c_1^2}{c_2} (R_{\text{pri}}^{\bar{x}})^2. \quad (19)$$

In practice, we collapse the constants into a single user parameter  $c$  and obtain the *ray footprint threshold*, which corresponds to the first term in Equation 5. Increasing  $c$  effectively shrinks the admissible footprint and enforces a more conservative notion of local smoothness, making reconnections rarer but more robust.

*Inverse ray footprint threshold.* The ray footprint threshold controls the change in area density at  $\mathbf{x}_k$ , but it does not account for how reconnection affects the BSDF sampling density  $p_k^{\bar{x}}(\omega_k)$  at the next vertex. This effect depends only on the change in outgoing direction  $-\omega_{k-1}$ , and can be severe for low-roughness glossy materials where the BSDF PDF is sharply peaked.

We therefore introduce the *inverse ray footprint*,

$$(p_k^{\bar{x}}(\omega_k)G(\mathbf{x}_k \rightarrow \mathbf{x}_{k-1}))^{-1}, \quad (20)$$

which we interpret as the footprint of a reverse-traced ray from  $\mathbf{x}_k$  towards  $\mathbf{x}_{k-1}$ . Assuming approximate reciprocity of the BSDF sampling PDFs for glossy materials,

$$p_k^{\bar{x}}(\omega_k) = p(\omega_k|\mathbf{x}_k, -\omega_{k-1}) \approx p(-\omega_{k-1}|\mathbf{x}_k, \omega_k), \quad (21)$$

this inverse footprint is proportional to the area density  $p(\mathbf{x}_{k-1}|\mathbf{x}_k, \omega_k)$  of the hit point of a reverse ray, via

$$p(-\omega_{k-1}|\mathbf{x}_k, \omega_k) = p(\mathbf{x}_{k-1}|\mathbf{x}_k, \omega_k)G(\mathbf{x}_k \rightarrow \mathbf{x}_{k-1}). \quad (22)$$

Reconnecting from  $\mathbf{y}_{k-1}$  instead of  $\mathbf{x}_{k-1}$  perturbs this density, and we can reuse the same footprint-smoothness argument as above: if

$$|\mathbf{y}_{k-1} - \mathbf{x}_{k-1}| < \sqrt{\frac{c_2}{p(\mathbf{x}_{k-1}|\mathbf{x}_k, \omega_k)}}, \quad (23)$$

then

$$\frac{|p(\mathbf{y}_{k-1}|\mathbf{x}_k, \omega_k) - p(\mathbf{x}_{k-1}|\mathbf{x}_k, \omega_k)|}{p(\mathbf{x}_{k-1}|\mathbf{x}_k, \omega_k)} < \epsilon. \quad (24)$$

Again using the geometric bound and assuming that the geometry term  $G(\mathbf{x}_k \rightarrow \mathbf{x}_{k-1}) \approx G(\mathbf{x}_k \rightarrow \mathbf{y}_{k-1})$  for distant reconnections, we arrive at a condition analogous to the ray footprint case:


$$(p_k^{\bar{x}}(\omega_k)G(\mathbf{x}_k \rightarrow \mathbf{x}_{k-1}))^{-1} > \frac{c_1^2}{c_2} (R_{\text{pri}}^{\bar{x}})^2. \quad (25)$$

### 3 CONDITIONAL VS. MARGINAL PDF IN FOOTPRINT THRESHOLD

To handle multi-layer material (e.g. material that contains both diffuse and specular components), ReSTIR PT uses lobe-specific connectability [Lin et al. 2022] – only the roughness of the sampled lobe is used to test against the roughness threshold. Correspondingly, the lobe index is preserved in reconnection. This yields a Jacobian determinant expressed in Equation 4. Since reconnection does not change the lobe type, it is reasonable to assume that the lobe sampling PMF remains nearly constant (as a result of shifting to a neighboring pixel that has similar material and similar viewing



**Fig. 1.** Replacing the ideal conditional PDF with marginal PDF for directional sampling results in shorter reconnection distances for glossy reflection for multi-lobed material as shown in this modified VEACH AJAR scene, where the reconnection distances shrinks as the diffuse component strength grows. This only results in slight quality loss.  $c = 0.02$  is used in this example.



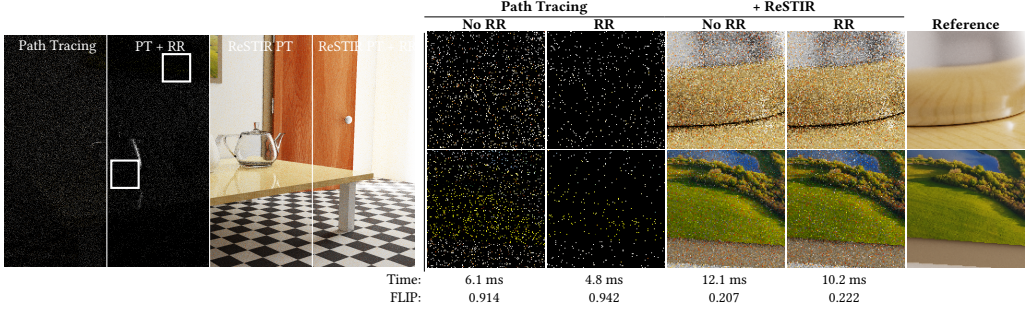
**Fig. 2.** Comparison between applying reconnection shift (left) and hybrid shift (right) on a teapot lid with a neural material that represents multiple material layers without explicit lobe or roughness information. Our hybrid shift uses the PDF-based threshold in place of the single-vertex roughness threshold to achieve a similar shift mapping effect, and provides significantly higher quality for the specular reflection.

direction) and focus on the change of the direction sampling PDF within the lobe when analyzing the threshold. Ideally, directional sampling PDFs  $p_{k-1}^x(\omega_{k-1}|\ell_{k-1})$  and  $p_k^x(\omega_k|\ell_k)$  conditioned on the lobes should be used instead of  $p_{k-1}^x(\omega_{k-1})$  and  $p_k^x(\omega_k)$ , respectively.

In practice, we simply use the marginal directional sampling PDFs (over all lobes) as we found that the quality difference is negligible. This results in the dual footprint threshold in Equation 5 we presented earlier. While this can sometimes cause premature reconnection on materials mixing diffuse and specular components (e.g. plastics), where the specular PDF tested against the threshold is reduced by the diffuse component, the effect is minor (Figure 1): a stronger reduction corresponds to a weaker specular contribution. The practical benefit is that we can reuse the same marginal PDF evaluation already needed for MIS weights, avoiding extra API complexity in the rendering system.

#### 4 PDF-BASED CONNECTIBILITY FOR BLACK-BOX MATERIAL

For general materials (BSDF), we do not always have a separation of lobes available, and roughness values might not be available. Still, we can modify our thresholds such that it can be used for any black-box material with only BSDF and PDF queries available. The footprint thresholds work as is with BSDF-sampling PDF substituted. To use the roughness threshold, we replace the roughness value with the reciprocal squared solid angle PDF that samples the  $\omega_{k-1}$  direction (inspired by



**Fig. 3.** Applying Russian roulette (RR) on path tracing yields much fewer non-zero pixels in the *VEACH AJAR* scene. However, with spatiotemporal reuse there is only slight noise increase.

spherical gaussians):

$$1/(p_{k-1}^x(\omega_{k-1}))^2 \geq \alpha_{\min}. \quad (26)$$

We find that using the same  $\alpha_{\min}$  tuned for the original roughness threshold yields satisfying results. Intuitively, a large PDF value suggests a glossy reflection which is likely to cause problem for reconnection. For multi-layer black-box material, large PDFs also suggest the "specular lobe" is being sampled. An example on neural material [Zeltner et al. 2024] is shown in Figure 2.

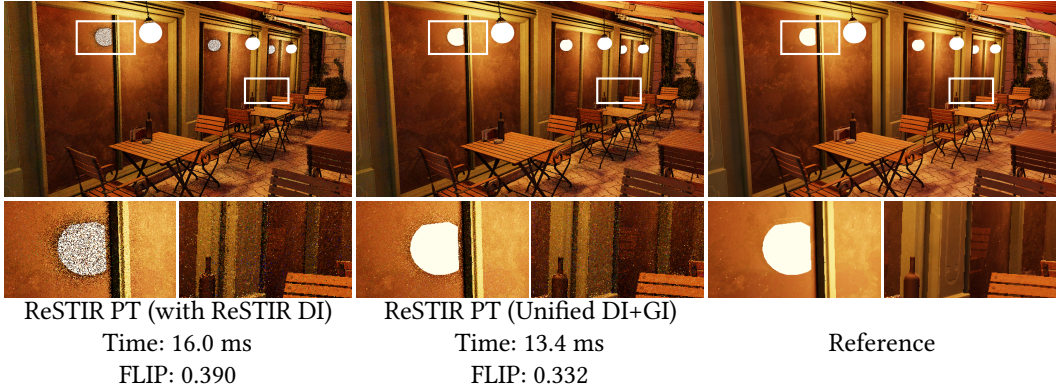
## 5 RIS-BASED NEE IN PRIMARY SPACE PARAMETERIZATION

*MIS Weights with RIS-Based NEE.* With  $M > 1$  ( $M$  being the candidate count) RIS for NEE, we replace the original single-sample MIS weight in the integrand to multi-sample MIS weight, i.e. for an NEE path  $\bar{x}$ ,  $\omega_1(\bar{x}) = Mp_1(\bar{x})/(Mp_1(\bar{x}) + p_2(\bar{x}))$  where 1 and 2 represents NEE and BSDF light sampling strategy, respectively.  $\omega_2(x)$  is defined similarly. Note that the actual NEE PDF is unknown due to the RIS process and we use the  $M$ -sample MIS formulation with the light sampling source PDF  $p_1$  assuming the RIS can approximate the quality of fully evaluating all  $M$  lights in NEE.


*PSS Implication.* We treat the RIS process as a black box and stick to our original primary sample space definition which assumes only one NEE sample is produced with a known PDF. Note that the actual sampling process does not have to be the same as the sampling process the PSS assumes. Therefore, the integrand (without lobe indexing) is still  $\omega_1(\bar{x})f(\bar{x})/p_1(\bar{x})$  for an NEE light sample, just with a modified MIS weight  $\omega_1$  compared to the single sample NEE. The effect of the RIS is incorporated into the UCW. Instead of 1 (assuming no Russian roulette), the UCW of the selected NEE sample becomes  $W_U = W_{\mathcal{X}(U)}^{\text{RIS}} p_1(\mathcal{X}(U))$  where  $U$  is the hypothetical random numbers that could have generate the same light sample in the original single-sample NEE sampling procedure and  $W_{\mathcal{X}(U)}^{\text{RIS}}$  is the path-space UCW produced by the RIS process, assuming area measure like in ReSTIR DI is used. The  $p_1(\mathcal{X}(U))$  term is also the Jacobian that converts between the parameterizations and we can verify that the original NEE sampling has  $W_{\mathcal{X}(U)} = 1/p_1(\mathcal{X}(U))$  which cancels the Jacobian to yield 1.

## 6 RUSSIAN ROULETTE IMPLEMENTATION DETAILS

The initial resampling can often be the most expensive stage in ReSTIR PT due to the code and data divergence in a path tracer. To reduce thread divergence in initial resampling, we apply Russian roulette to reduce the average path length and also lower the chance of excessively long paths. Usually, the primary sample space (PSS) is extended with random numbers for Russian roulette, which can create caveats for shift mapping: naively replaying the random numbers to perform Russian roulette tests in hybrid shift can increase shift failure, as a survived path can sometimes be shifted to a killed path, invalidating the sample.



**Fig. 4.** By including direct illumination in sampling and disabling ReSTIR DI [Bitterli et al. 2020], the Unified version significantly improves the DI sampling, especially for specular reflection, thanks to hybrid shift and pairwise MIS. Using light tiles for NEE also improved lighting quality of subsequent bounces (reflection of the chair in the second inset). Frame time also reduces due to eliminating a separate ReSTIR pass.



**Fig. 5.** Evaluation of our color noise reduction method in ZERO DAY.  $M = 4$  for spatial resampling (3 spatial neighbors) in this case. By accumulating the vectorized form of resampling weight in spatial reuse, we achieve significant color noise reduction almost for free.

Our approach is to deterministically select survival when shifting a path. This would require extra Jacobian calculation in a PSS extended with Russian roulette dimensions. We circumvent this complexity by decoupling the definition of primary sample space from the actual sampling process – the paths are defined in the original primary sample space without Russian roulette and the Russian roulette only alters the initial sample’s PDF. In practice, we change the uniform  $p(\bar{\mathbf{u}}) = 1$  to  $p(\bar{\mathbf{u}}) = \prod_{i=1}^{d-1} q_i(\bar{\mathbf{u}})$  where  $q_i(\bar{\mathbf{u}})$  are survival probabilities at each bounce. This way, although Russian roulette can make the initial samples noisier, the ReSTIR PT output quality is affected to a much smaller extent thanks to spatiotemporal reuse (Figure 3).

## 7 INDIVIDUAL EVALUATIONS ON DI-GI UNIFICATION, COLOR NOISE AND DISOCCLUSION NOISE REDUCTION

Figure 4, Figure 5, and Figure 6 present evaluations of our DI-GI unification, color noise reduction, and disocclusion noise reduction methods, respectively.



**Fig. 6.** Applying dual motion vectors [Zeng et al. 2021] provides a dramatic reduction of disocclusion noise which is severe in the VEACH AJAR due to undersampling. The result contains some correlated noise but does not show "copy-pasting" artifacts.

## 8 RESERVOIR COMPRESSION

Reducing reservoir sizes helps reduce the bandwidth cost inside the reuse passes where each pixel reads its own reservoir from the memory and randomly access the reservoirs in a local neighborhood. Lin et al. [2022] uses 88-byte reservoirs for ReSTIR PT where attributes related to the reconnection vertex occupy a lot of the reservoir size. We found that many of the attributes can be compressed in a lossy format without affecting the quality. We list the details in Algorithm 1 where we shrink the reservoir size to 64 bytes.

**Algorithm 1** A compression of the reservoir struct shrinks its size from 88 bytes to 64 bytes.

```

struct Reservoir {
  float W; float3 F; // F is the integrand, pHat = luminance(F).
  uint initRandomSeed, rcVertexRandomSeed;
  float M; uint pathFlags; // Pack M as 8 bit integer into pathflags.
  uint rcVertexInstanceID, rcVertexPrimitiveIndex;
  float2 → uint rcVertexBarycentrics; // Compress as two 16-bit Unorms.
  float3 → uint rcVertexWi; // Octahedron encoding as two 16-bit Unorms.
  float3 rcVertexRadiance;
  float4 → float2 rcVertexCachedValues; // Combine the separately stored
  Jacobian terms  $p(\omega_k), G(x_{k-1}, x_k), p(\omega_{k+1})$  (float3) of the base path into their
  product (float); NEE Light PDF (float) which is separately used to compute
  path MIS weight is unchanged.
}

```





## 9 COMPLETE PERFORMANCE RESULTS

Table 1 provides complete performance results of the four tested scenes at the default 1920×1080 resolution. Table 2 studies performance scaling across 540p, 1080p, and 2160p render resolutions. Compared with the baseline ReSTIR PT, our ReSTIR PT Enhanced has a smaller cost increase as resolution rises. In particular, our frame time increases by 4.03× from 1080p to 2160p (a 4× increase of pixels), whereas the baseline’s frame time increases by 4.57×. This is likely because our optimizations reduce pressure on the memory system, leading to better scaling at higher resolutions.

## 10 ABLATION OF FOOTPRINT THRESHOLD PARAMETER CHOICE

A parameter sweep for  $c$  for the dual ray footprint threshold is provided in Figure 8. The optimal  $c$  are in the range of 0.005-0.04 across tested scenes with a wide variety of scales and materials, and

**Table 1.** Performance Results: Frame Time and Pass Times (in milliseconds)

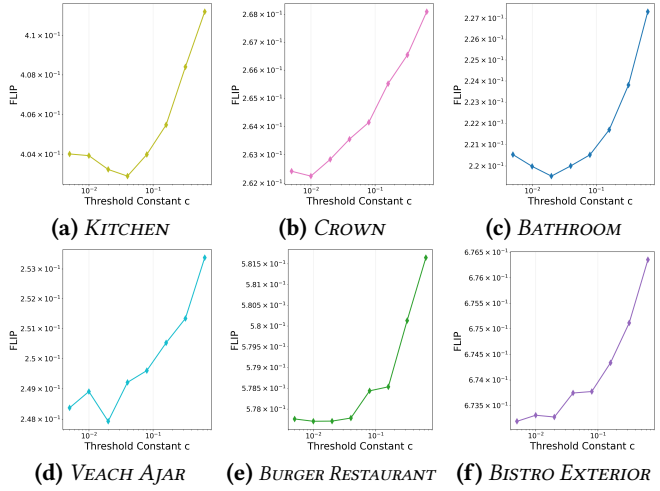
Scene	Method	Frame Time	Initial Sampling	Temporal Reuse	Spatial Reuse	ReSTIR DI and others
	Baseline	21.99	3.57	2.99	12.23	3.20
	+Code microop	19.34	3.54	2.19	10.44	3.17
	+Forced NEE reconnect	15.75	3.98	1.93	6.89	2.95
	+Replay compaction	13.02	4.16	1.40	4.50	2.96
	+Paired spatial reuse	12.10	4.56	1.53	2.88	3.13
	+Russian roulette	10.20	3.17	1.34	2.67	3.02
	+Unify DI & GI	7.64	2.89	1.29	2.71	0.75
	+New thresholds	8.74	3.09	1.72	2.92	1.01
	+All improvements	9.89	3.14	1.81	3.17	1.77
	Baseline	46.76	17.74	6.88	15.92	6.22
	+Code microop	43.82	17.74	5.70	14.25	6.13
	+Forced NEE reconnect	40.55	18.02	5.00	11.45	6.08
	+Replay compaction	36.51	17.17	4.29	8.81	6.24
	+Paired spatial reuse	32.85	17.35	4.36	4.81	6.33
	+Russian roulette	20.16	6.29	3.07	4.50	6.30
	+Unify DI & GI	17.47	8.98	3.10	3.93	1.46
	+New thresholds	19.43	10.01	3.62	4.23	1.57
	+All improvements	20.70	10.12	4.17	4.33	2.08
	Baseline	49.29	15.69	7.25	19.10	7.25
	+Code Microop	46.87	16.25	5.82	17.17	7.63
	+Forced NEE reconnect	43.61	17.78	4.45	13.21	8.17
	+Replay compaction	40.30	19.35	3.10	9.10	8.75
	+Paired spatial reuse	38.63	21.13	3.31	5.11	9.08
	+Russian roulette	22.14	6.30	2.78	4.78	8.28
	+Unify DI & GI	16.69	9.33	2.48	3.88	1.00
	+New thresholds	18.67	8.81	4.30	3.97	1.59
	+All improvements	19.54	8.96	4.67	4.07	1.84
	Baseline	24.87	5.34	4.09	11.91	3.53
	+Code microop	21.90	5.26	2.94	10.37	3.33
	+Forced NEE reconnect	19.10	5.99	2.20	7.36	3.55
	+Replay compaction	17.41	6.50	1.84	5.55	3.52
	+Paired spatial reuse	16.49	7.21	1.99	3.63	3.66
	+Russian roulette	13.59	5.06	1.78	3.38	3.37
	+Unify DI & GI	10.35	4.69	1.68	3.18	0.80
	+New thresholds	11.20	4.80	2.04	3.27	1.09
	+All improvements	11.99	4.84	2.15	3.35	1.65
Average	Baseline	35.73	10.59	5.30	14.79	5.05
	+Code microop	32.98	10.70	4.16	13.06	5.07
	+Forced NEE reconnect	29.75	11.44	3.40	9.73	5.19
	+Replay compaction	26.81	11.79	2.66	6.99	5.37
	+Paired spatial reuse	25.02	12.56	2.80	4.11	5.55
	+Russian roulette	16.52	5.21	2.24	3.83	5.24
	+Unify DI & GI	13.04	6.47	2.14	3.43	1.00
	+New thresholds	14.51	6.68	2.92	3.60	1.31
	+All improvements	15.53	6.77	3.20	3.73	1.83

**Table 2.** Performance vs. Resolution. "Ours" includes all our optimizations and quality improvements. Frame time (milliseconds) is reported for each scene and resolution combination.

Method	Resolution	VEACH AJAR	CAROUSEL	OPERA HOUSE	BATHROOM	Avg. time scaling
Baseline	960×540	5.75	14.28	13.34	6.53	
	1920×1080	21.99	46.76	49.29	24.87	3.65×
	3840×2160	97.66	208.63	234.12	114.63	4.57×
Ours	960×540	2.70	6.66	5.26	3.16	
	1920×1080	15.53	20.70	19.54	11.99	3.57×
	3840×2160	62.34	77.99	82.10	50.13	4.03×



Fig. 7. Scenes



**Fig. 8.** Ablation of constant  $c$  values for the dual footprint threshold using the FLIP metric, testing  $c = 0.005$ – $0.64$  across six scenes with very different scene dimensions and materials. KITCHEN, CROWN, BATHROOM, VEACH AJAR, BURGER RESTAURANT, BISTRO EXTERIOR achieve minimum FLIP at  $c = 0.04, 0.01, 0.02, 0.02, 0.01, 0.005$ , respectively. This gives a  $c = 0.0175$  scene average, proving the near-optimal choice of  $c = 0.02$ .




**Fig. 9.**  $c = 0.02$  and  $c = 0.005$  do not have obvious difference for the dual footprint threshold in BURGER RESTAURANT (no correlation reduction technique used in the result).

the cross-scene average optimal  $c = 0.0175 \approx 0.02$ . We can see that diffuse scenes (e.g. Figure 9) are less sensitive to the  $c$  choice than glossy scenes (e.g. Figure 10).


## 11 ADDITIONAL ABLATION STUDIES OF OUR DUAL FOOTPRINT THRESHOLD

To further validate the effectiveness of our method, we also compare against two alternative methods to define adaptive thresholds that are directly derived from prior work.


The first alternative is directly using the path footprint threshold as in Neural Radiance Cache [Müller et al. 2021]. While this is a good approach for determining when to stop path tracing to read from the radiance cache, it is insufficient to function as a robust reconnection threshold, mainly because it does not consider the BSDF of the reconnection vertex. Figure 11 shows a case where the path footprint threshold allows reconnection vertex to be on a delta surface, causing poor caustics quality.



**Fig. 10.**  $c = 0.02$  is significantly better than  $c = 0.005$  for the dual footprint threshold in KITCHEN (no correlation reduction technique used in the result).



**Fig. 11.** Ignoring the reconnection vertex material, the path footprint threshold causes invalid reconstructions on the glass egg and noisy caustics ( $c = 0.02$ ) in regions farther from the refractor. Achieving similar caustic quality as ours needs at least  $c = 2$ , which degrades diffuse reflection quality. This means no  $c$  value for the path footprint threshold can match our quality in this case.



**Fig. 12.** As can be shown in this case in KITCHEN scene, having two BSDF-sampling PDFs in the threshold test rejects reconnection between glossy surfaces over-aggressively ( $c = 0.02$  right inset). When the thresholding constant  $c$  is lowered to 0.0002, the glossy reflection quality is still a bit worse than our method, meanwhile the noise at the diffuse wall corner (left inset) starts to increase. This means no  $c$  value for Jacobian bounding can match our quality in this case.

The second alternative is replacing the dual footprint threshold with a single "Jacobian bounding" threshold. Basically Equation 5 is changed into

$$(p_{k-1}^x(\omega_{k-1})G(\mathbf{x}_{k-1} \leftrightarrow \mathbf{x}_k)p_k^x(\omega_k))^{-1} \geq \frac{c}{100} \frac{\|\mathbf{x}_0 - \mathbf{x}_1\|^2}{\langle n_{\mathbf{x}_1}, \widehat{\mathbf{x}_1\mathbf{x}_0} \rangle / (8\pi^2)}, \quad (27)$$


where the RHS is that of Equation 5 multiplied by  $2\pi$  (the reciprocal PDF of sampling a hemisphere uniformly) to match the threshold scales after including an additional PDF value in LHS. Intuitively, this bounds the PSS-to-area Jacobian determinant, so that the reconnection is less likely to have extreme Jacobians. This method also imitates clamping the spike-inducing  $\rho(\mathbf{x}_0 \rightarrow \mathbf{x}_1 \rightarrow \mathbf{x}_2)G(\mathbf{x}_1 \leftrightarrow \mathbf{x}_2)\rho(\mathbf{x}_1 \rightarrow \mathbf{x}_2 \rightarrow \mathbf{x}_3)$  term in VPL contribution [Davidovič et al. 2010; Hašan et al. 2009]. As can be seen in Figure 12, the Jacobian-bounding method is suboptimal compared to our method as it discourages reconnection between two glossy surfaces too strongly.

## 12 BASELINE COMPARISON WITH AND WITHOUT DECORRELATION

Figure 13 provides a three-way comparison between the baseline, our method without the decorrelation technique (Section 5 of the main paper), and our default method with decorrelation enabled. Compared to the baseline, both versions of our method provides the same level of speedup and enhancement of image quality (when correlation is ignored) due to various algorithmic improvements. Our decorrelation technique further suppresses correlated noise at a slight performance cost.

## REFERENCES

- 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 Transactions on Graphics (Proceedings of SIGGRAPH)* 39, 4 (July 2020).
- Tomáš Davidovič, Jaroslav Krivánek, Miloš Hašan, Philipp Slusallek, and Kavita Bala. 2010. Combining global and local virtual lights for detailed glossy illumination. *ACM Transactions on Graphics (TOG)* 29, 6 (2010), 1–8.
- Miloš Hašan, Jaroslav Krivánek, Bruce Walter, and Kavita Bala. 2009. Virtual spherical lights for many-light rendering of glossy scenes. In *ACM SIGGRAPH Asia 2009 papers*. 1–6.
- Daqi Lin, Markus Kettunen, Benedikt Bitterli, Jacopo Pantaleoni, Cem Yuksel, and Chris Wyman. 2022. Generalized Resampled Importance Sampling: Foundations of ReSTIR. 41, 4 (2022), 75:1–75:23. <https://doi.org/10.1145/3528223.3530158>
- Thomas Müller, Fabrice Rousselle, Jan Novák, and Alexander Keller. 2021. Real-time neural radiance caching for path tracing. *arXiv preprint arXiv:2106.12372* (2021).
- Tizian Zeltner, Fabrice Rousselle, Andrea Weidlich, Petrik Clarberg, Jan Novák, Benedikt Bitterli, Alex Evans, Tomáš Davidovič, Simon Kallweit, and Aaron Lefohn. 2024. Real-time neural appearance models. *ACM Transactions on Graphics* 43, 3 (2024), 1–17.
- Zheng Zeng, Shiqiu Liu, Jinglei Yang, Lu Wang, and Ling-Qi Yan. 2021. Temporally Reliable Motion Vectors for Real-time Ray Tracing. In *Computer Graphics Forum*, Vol. 40. Wiley Online Library, 79–90.



**Fig. 13.** Comparison between original ReSTIR PT [Lin et al. 2022], our method with no decorrelation, and our default method. FLIP error increases slightly because FLIP is relatively insensitive to correlation artifacts, but sensitive to luminance differences introduced by bias.