

Conditional Resampled Importance Sampling and ReSTIR

Markus Kettunen*
NVIDIA
Finland
mkettunen@nvidia.com

Daqi Lin*
NVIDIA
USA
daqil@nvidia.com

Ravi Ramamoorthi
NVIDIA and UC San Diego
USA
ravir@cs.ucsd.edu

Thomas Bashford-Rogers
University of Warwick
UK
thomas.bashford-
rogers@warwick.ac.uk

Chris Wyman
NVIDIA
USA
chris.wyman@acm.org




Figure 1: Our new conditional RIS theory enables new types of unbiased subpath reuse by resampling in conditional probability spaces. To show the theory has practical use, we prototype an algorithm resampling multiple ReSTIR-driven path suffixes in a photon map like final gather. While our proof-of-concept is unoptimized, we compare with two state-of-the-art methods without conditional resampling, including a final gather using West et al.’s [2022] marginal multiple importance sampling (MMIS) and full-path resampling using Lin et al.’s [2022] ReSTIR PT sample code. The TOWER BRIDGE [Pobursky 2021] is lit by the SHANGHAI BUND probe; the camera sees an almost entirely indirectly lit region. ReSTIR PT is very fast, but complex lighting plus specular surfaces can cause large spatiotemporal correlations, boiling, and color shifts (left inset). While currently more expensive, our subpath resampling gives spatiotemporally stable results without visible correlations. Compared to an MMIS gather, our prototype improves quality given a similar ray budget. (Bottom) Below each image we show (x, t) plots taken from videos (without movement); rows come from sequential video frames, so temporal correlations appear as vertical blobs and spatial correlations show up as horizontal blobs. All techniques are unbiased, converging to reference in time, but results here use only one full path per-pixel for integration.

ABSTRACT

Recent work on generalized resampled importance sampling (GRIS) enables importance-sampled Monte Carlo integration with random variable weights replacing the usual division by probability density. This enables very flexible spatiotemporal sample reuse, even if neighboring samples (e.g., light paths) have intractable probability densities. Unlike typical Monte Carlo integration, which samples according to some PDF, GRIS instead resamples existing samples. But resampling with GRIS assumes samples have tractable marginal

contribution weights, which is problematic if reusing, for example, light *subpaths* from unidirectionally-sampled paths. Reusing such subpaths requires conditioning by (non-reused) segments of the path prefixes.

In this paper, we extend GRIS to conditional probability spaces, showing correctness given certain conditional independence between integration variables and their unbiased contribution weights. We show proper conditioning when using GRIS over randomized conditional domains and how to formulate a joint unbiased contribution weight for unbiased integration.

To show our theory has practical impact, we prototype a modified ReSTIR PT with a final gather pass. This reuses subpaths, postponing reuse at least one bounce along each light path. As in photon mapping, such a final gather reduces blotchy artifacts from sample correlation and reduced correlation improves the behavior of modern denoisers on ReSTIR PT signals.

*Joint first authors; equal contribution.

Permission to make digital or hard copies of part or all 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 third-party components of this work must be honored. For all other uses, contact the owner/author(s).

SA Conference Papers '23, December 12–15, 2023, Sydney, NSW, Australia

© 2023 Copyright held by the owner/author(s).

ACM ISBN 979-8-4007-0315-7/23/12.

<https://doi.org/10.1145/3610548.3618245>

CCS CONCEPTS

• **Computing methodologies** → **Ray tracing**; • **Mathematics of computing** → **Resampling methods**.

ACM Reference Format:

Markus Kettunen, Daqi Lin, Ravi Ramamoorthi, Thomas Bashford-Rogers, and Chris Wyman. 2023. Conditional Resampled Importance Sampling and ReSTIR. In *SIGGRAPH Asia 2023 Conference Papers (SA Conference Papers '23)*, December 12–15, 2023, Sydney, NSW, Australia. ACM, New York, NY, USA, 11 pages. <https://doi.org/10.1145/3610548.3618245>

1 INTRODUCTION

Various recent theoretical advancements have improved the state of the art in Monte Carlo integration significantly, particularly for light transport in real-time rendering. Generalizations of resampled importance sampling (RIS) [Talbot et al. 2005] enable unbiased sample reuse between complicated integration domains, leading to reservoir-based spatiotemporal importance resampling (ReSTIR) [Bitterli et al. 2020], which promises to amortize path tracing costs over large numbers of pixels, both within and across frames.

A key aspect of ReSTIR’s speed and unbiasedness is careful multiple importance sampling (MIS) weight computation [Veach and Guibas 1995b], as each spatiotemporal neighbor acts as a unique sampler. If reusing samples from differing domains, correct MIS weights are vital to harmonizing supports to avoid under- and over-counting different regions of the integration domain.

West et al. [2022] extended MIS to marginal PDFs (MMIS), where probability densities are not pointwise evaluable. This seems vital in generalized resampled importance sampling (GRIS) [Lin et al. 2022], which replaces the notion of a PDF with *unbiased contribution weights* (UCW) that are, by definition, not pointwise evaluable. But GRIS theory does not handle *conditional* unbiased contribution weights, limiting the sample reuse available to ReSTIR.

In this paper, we extend the generalized resampled importance sampling theory to handle conditional probability spaces. This enables (iteratively) driving conditioning and conditioned variables using previous RIS- or ReSTIR-based samples. This was motivated by a practical problem: in ReSTIR PT [Lin et al. 2022], challenging light paths’ next event estimation (NEE) queries are often occluded. We wondered if we could replace these single-segment NEE queries with important, multi-segment subpaths reused from neighbors.

To demonstrate our theory, we add a final gather to ReSTIR PT. We shoot paths unidirectionally from the camera, update them via ReSTIR, but only reuse a *portion*, discarding (at least) one segment from the start before reuse. This requires conditional ReSTIR; subpaths are conditioned by the (discarded) prefixes initially used to generate the paths. The algorithm remains unbiased.

Our specific theoretical contributions include:

- Extending the definition of unbiased contribution weights to allow conditional and joint UCWs.
- Demonstrating how to do conditional RIS and ReSTIR.
- A general MIS scheme for integrating by paths reused with conditional shift mappings—without knowing PDFs.
- A framework for spatiotemporal improvement and reuse of unidirectionally sampled suffix paths.

- Finding similarity between ReSTIR PT and photon mapping algorithms, allowing a final gather for reducing correlations.

This finishes the evolution begun by Lin et al. [2022] to allow replacing PDFs in Monte Carlo integration ($f(X)/p(X)$) with unbiased contribution weights ($f(X)W_X$). Now, UCWs can exist in conditional probability spaces, $W_{X|Z}$, depend jointly on multiple variates, and be marginalized with respect to specific random variables.

2 KEY CHALLENGE IN RIS AND RESTIR

Before exploring prior work and our new theory, let’s review the key difficulty for sample reuse in resampling (e.g., Talbot et al. [2005] and follow ups). In theory, sample reuse is great; the question is “can we?” Considering the high-dimensionality, it is not clear unrelated paths can help integrate over the same domain. Path dimensionality may vary and it is unclear if path measures match; arbitrary sample reuse can feel a bit like asking, “how many grams are in a meter?”

Thus, the work designing ReSTIR [Bitterli et al. 2020; Lin et al. 2022] focused on a few key issues:

- Knowing the sampling domains,
- Matching domains between samples for reuse,
- Shifting samples’ domains to enable more reuse,
- Ensuring we sample the full integration domain, and
- Avoiding double counting parts of the domain.

Discarding sampling domain data is one reason why post-process denoising is fundamentally biased; how can one ensure unbiasedness when averaging final pixel colors that estimate different integrands?

We extend ReSTIR, asking “can we reuse *part* of a path?” This means readdressing these issues, as subpath reuse makes the domains conditional. For example, if reusing from the 4th path vertex, that subpath was picked relative to the now unused 3rd vertex.

A key difference between our work and prior methods for sampling conditional spaces (e.g., West et al. [2022]) is that we enable a streaming compute model, where reservoirs store individual samples representing aggregations of many (sub-)paths.

3 PREVIOUS WORK AND BACKGROUND

Prior research explores methods for both *path reuse* and *path filtering*. Here, we define path reuse as reusing entire samples. Simple methods copy all neighbor path vertices to the current pixel [Bekaert et al. 2002; Ouyang et al. 2021]. Gradient domain rendering added *shift mappings*, allowing more complex sample reuse between integration domains [Bauszat et al. 2017; Hua et al. 2019]. Similar reuse ideas arise in Metropolis techniques [Veach and Guibas 1997].

Path filtering usually suggests a biased smoothing, typically at path vertices, that connects one path’s *prefix* to other path *suffixes* [Binder et al. 2019; Keller et al. 2014; West et al. 2020]. West et al. [2022] noted path filtering can be unbiased, albeit cost prohibitive due to many connections needed in the filtering kernel. We show suffix resampling yields high quality while borrowing just one suffix.

Generally, our theory fits into this reuse and filtering space but allows *conditional* reuse of subpaths (similar to Tessari et al. [2017] or West et al. [2022]) in the context of resampled importance sampling [Lin et al. 2022], where reused subpaths come from large sample aggregations (reservoirs) from spatiotemporal neighbors.

To demonstrate practical benefits of our theory, we build a proof-of-concept final gather to connect multiple path prefixes to one or more suffixes. Such gathering is not new, per se, being used in photon mapping [Jensen 2001] and path filtering [Binder et al. 2019]. Deng et al. [2021] and West et al. [2022] also showed iterative gathering (or multi-vertex filtering) to be valuable; our theory may enable such ideas in the context of real-time GPU-accelerated ReSTIR variants.

Interestingly, conditional subpath reuse resembles bidirectional path tracing [Lafortune and Willems 1993; Veach and Guibas 1995a] without sampling from lights. Camera paths find key lights; subpaths hitting a light are reused for neighbor pixels as light subpaths. Another view of our prototype is as an unbiased radiance cache, in contrast to prior work that adds bias by estimating density of virtual point lights (traced from the camera) [Segovia et al. 2006] or applying ReSTIR to reuse biased light probes [Majercik et al. 2021].

3.1 Multiple Importance Sampling

Many learn multiple importance sampling (MIS) through the lens of Veach and Guibas’s [1995b] example joining light and BSDF samples into an improved estimator. Crucially, MIS allows weighing samples reused across varying integrands and integration domains. Elvira et al. [2019] and Deng et al. [2019] pushed such reuse, ultimately expanding MIS [West et al. 2020] to enable randomly selected subsets of uncountable continuums of sampling techniques.

Many MIS methods require readily evaluable PDFs for each estimator used. Such PDFs are not available when resampling [Talbot et al. 2005]; Lin et al. [2022] showed MIS weights can be computed using RIS’s target functions that only *approximate* estimator PDFs.

West et al. [2022] introduced marginal MIS, allowing combining sampling methods whose probability distributions are known only conditionally, given some random variables. This enables MIS between conditional PDFs. In this paper we further allow combining and chaining such conditional distributions using RIS and ReSTIR.

3.2 RIS and ReSTIR

Resampled importance sampling (RIS) [Talbot et al. 2005] and reservoir spatiotemporal importance resampling (ReSTIR) [Bitterli et al. 2020] propose that neighbor sample distributions likely approximate the current integrand better than any analytic distribution. With iterative spatiotemporal bootstrapping, ReSTIR basically draws samples from *approximately perfect* distributions (see Devroye [1986]).

Similar ideas appear elsewhere, e.g. Metropolis [Veach and Guibas 1997] or simpler path reuse methods [Bekaert et al. 2002], but ReSTIR dramatically accelerates this with streaming computation via weighted reservoir sampling [Chao 1982]. Essentially, ReSTIR allows massive sample amortization, spreading cost over many pixels, without correspondingly higher storage costs.

This builds on the RIS estimator [Talbot et al. 2005], where f is our integrand (typically the path contribution function):

$$\langle I \rangle_{\text{ris}} = \frac{1}{N} \sum_{i=1}^N \left[\frac{f(Y_i)}{\hat{p}(Y_i)} \frac{1}{M} \sum_{j=1}^M \frac{\hat{p}(X_{ij})}{p(X_{ij})} \right]. \quad (1)$$

Here, each of the N samples Y_i is resampled from M independent candidates X_{ij} from some distribution with PDF p . The X_{ij} are

reweighted by some *target function* \hat{p} , and the Y_i are drawn from the X_{ij} proportional to their new weights. These N samples then estimate f via Monte Carlo importance sampling.

Standard Monte Carlo estimators sum over samples of the form $f(Y_i)/p(Y_i)$, but the RIS estimator sums samples $f(Y_i)W_{Y_i}$, where

$$W_Y = \frac{1}{\hat{p}(Y)} \frac{1}{M} \sum_{j=1}^M \frac{\hat{p}(X_j)}{p(X_j)}. \quad (2)$$

Equation 2 includes another Monte Carlo estimator $\hat{p}(X_j)/p(X_j)$. Replacing it by RIS estimator $\hat{p}(X_j)W_{X_j}$ allows chaining RIS estimates, which is the key idea behind ReSTIR [Bitterli et al. 2020].

Sample reuse in RIS that leads to these weights W is one example of a more general idea, that of *unbiased contribution weights* introduced by Lin et al. [2022]:

Definition 3.1. An *unbiased contribution weight* (or UCW) for random variable X , is any real-valued random variable W_X for which

$$\mathbb{E}[f(X)W_X] = \int_{\text{supp}(X)} f(x) dx, \quad (3)$$

for any integrable function $f : \Omega \rightarrow \mathbb{R}$.

Note that throughout the paper, we denote random variables with capitals X and ordinary variables with lower case x .

Unlike traditional Monte Carlo sampling, where samples must be drawn from tractable PDFs p , UCWs allow use of more complex distributions. Regular PDFs act as UCWs using $W_X = 1/p_X(X)$, but any random variable can be a UCW as long as $\mathbb{E}[W_X|X] = 1/p_X(X)$.

While standard PDFs p can also come in conditional, marginal, or joint forms, prior ReSTIR work did not enable UCWs with such forms [Bitterli et al. 2020; Lin et al. 2022, 2021]. We extend UCWs to conditional probability spaces, allowing resampling in those domains. Such needs arise in complex rendering algorithms, e.g., if one wants to reuse only *part* of a previously sampled light path.

4 CONDITIONAL UCWS

A conditional PDF $p_{X|Y}$ can be seen as the PDF of X in a conditional probability space where random variable Y receives a specified value. Here Y is constant, so $f/p_{X|Y}$ estimates the conditional expectation $\mathbb{E}[f/p_{X|Y}|Y]$. But what if we only have an unbiased estimate of $1/p_{X|Y}$, e.g., by conditional RIS (Section 5)?

The estimator $f/p_{X|Y}$ has the conditional expectation

$$\mathbb{E} \left[\frac{f(X)}{p_{X|Y}(X|Y)} \middle| Y \right] = \int_{\text{supp}(X|Y)} f(x) dx, \quad (4)$$

where $\text{supp}(X|Y)$ contains values X possible with positive PDF given Y , $p_{X|Y}(X|Y) > 0$. Here f may depend on Y , as Y is fixed.

This $\mathbb{E}[\cdot|Y]$ can be interpreted as a traditional expectation in the conditional probability space where Y has fixed value. We observe that applying UCWs in such a conditional space naturally leads to a definition of *conditional unbiased contribution weights*.

Definition 4.1. A conditional unbiased contribution weight $W_{X|Y}$ for random variable X , given Y , is any real-valued random variable $W_{X|Y}$ for which

$$\mathbb{E}[f(X)W_{X|Y}|Y] = \int_{\text{supp}(X|Y)} f(x) dx, \quad (5)$$

for any integrable function $f: \Omega \rightarrow \mathbb{R}$.

Similar to traditional UCWs, an immediate follow-up is that $W_{X|Y}$ has conditional expectation $\mathbb{E}[W_{X|Y} | X, Y] = 1/p_{X|Y}(X|Y)$.

4.1 Joint UCWs

Knowing a marginal PDF p_Y and a conditional PDF $p_{X|Y}$ allows unbiased integration with pairs (X, Y) : the product $p_Y(y) p_{X|Y}(x|y)$ yields the joint PDF $p_{X,Y}(x, y)$. Does knowing W_Y and $W_{X|Y}$ allow unbiased integration? Let us approach this with an example.

We integrate $f(x, y)$ over the unit square with points (X_1, X_2) , with X_2 sampled conditionally on X_1 . Multiplying f by $W_{X_2|X_1}$ estimates the conditional expectation

$$\mathbb{E}[f(X_1, X_2)W_{X_2|X_1}|X_1] = \int_{\text{supp}(X_2|X_1)} f(X_1, x_2) dx_2, \quad (6)$$

which integrates over X_2 for the fixed X_1 . The right-hand-side expression only depends on X_1 , so now multiplying by W_{X_1} yields the integral over the full square:

$$\mathbb{E}[\mathbb{E}[f(X_1, X_2)W_{X_2|X_1}|X_1] W_{X_1}] = \int_{\text{supp}(X_1)} \int_{\text{supp}(X_2|X_1)} f(x_1, x_2) dx_2 dx_1. \quad (7)$$

Is it then true that whenever we have random variables X_1 and X_2 with UCWs W_{X_1} and $W_{X_2|X_1}$, then $f(X_1, X_2) W_{X_1} W_{X_2|X_1}$ unbiasedly estimates the integral of f over the joint support of (X_1, X_2) , making $W_{X_1} W_{X_2|X_1}$ a joint UCW for (X_1, X_2) ? The answer is *no*.

Without care, this key subtlety can lead to bias (see the supplemental document [Section S.1.2](#) for an example). The expectation of estimator $f(X_1, X_2)W_{X_1}W_{X_2|X_1}$ can be written

$$\mathbb{E}[\mathbb{E}[f(X_1, X_2)W_{X_1}W_{X_2|X_1}|X_1]], \quad (8)$$

which differs from [Equation 7](#). To yield the correct expectation, W_{X_1} must move to the outer expectation. This requires W_{X_1} be *conditionally independent* of $f(X_1, X_2)W_{X_2|X_1}$ given X_1 , i.e., the expressions must not depend on the same random variables, except for X_1 . This can be achieved by the following rule:

THEOREM 4.1. If X_2 and $W_{X_2|X_1}$ are conditionally independent of W_{X_1} , given X_1 , then

$$W_{X_1, X_2} = W_{X_1} W_{X_2|X_1} \quad (9)$$

is a joint unbiased contribution weight for $X = (X_1, X_2)$.

If X_2 and $W_{X_2|X_1}$ share dependencies with W_{X_1} , besides X_1 , then UCWs W_{X_1} and $W_{X_2|X_1}$ must be conditional on them, i.e., integration must succeed if we treat shared random variables as constant.

5 CONDITIONAL RIS AND INTEGRATION

As in [Talbot et al. \[2005\]](#) and [Lin et al. \[2022\]](#), our new conditional unbiased contribution weights can be computed with RIS. In this case, reused samples' domains often will not cover the full support of integrand f . Consider random variable X_1 and some dependent random variables arranged into a vector Z , and X_1 's conditional UCW $W_{X_1|Z}$ (for example, X_1 could be a path suffix and Z could be its prefix). By definition, this UCW can integrate any function f , potentially dependant on Z , in the conditional support $\text{supp}(X_1|Z)$:

$$\mathbb{E}[f(X_1)W_{X_1|Z}|Z] = \int_{\text{supp}(X_1|Z)} f(x_1) dx_1. \quad (10)$$

But if $\text{supp}(X_1|Z)$ does not cover f 's support, our estimate is biased. Ensuring unbiasedness requires at least one sample covering otherwise uncovered regions where $f(x) \neq 0$. We add one *canonical* sample X_2 (see [Lin et al. \[2022\]](#)), with UCW $W_{X_2|Z}$ and known to cover f 's entire support; the sampling procedure may depend on Z .

Combining $f(X_1)W_{X_1|Z}$ and $f(X_2)W_{X_2|Z}$ gives an unbiased estimate if we appropriately pick MIS weights. If samples were in one domain with known PDFs, we could use the balance heuristic,

$$m_i(x|Z) = \frac{p_i(x|Z)}{p_1(x|Z) + p_2(x|Z)}, \quad (11)$$

to give the unbiased estimate

$$m_1(X_1|Z) f(X_1)W_{X_1|Z} + m_2(X_2|Z) f(X_2)W_{X_2|Z}. \quad (12)$$

The conditional notation $m_i(\cdot|Z)$ has no deeper meaning: we could denote $m_{Z,i}(x)$ as a parametrized function family (e.g., [West et al. \[2020, 2022\]](#)). For future brevity, we occasionally implicitly drop this dependency on Z . (In fact, we had already dropped it from f).

In practice, samples X_i often arise via resampling, making an unmodified balance heuristic unusable as we use UCWs with unknown PDFs p_i . [Lin et al. \[2022\]](#) used ReSTIR target functions \hat{p}_i as PDF proxies for MIS, enabling sample reuse between domains via shift mappings and their Jacobians. Below, we borrow and expand on [Lin et al. \[2022\]](#) by interpreting their expressions with an implicit Z dependency and in a conditional probability space.

5.1 Resampling

Above, we enabled unbiased integration in conditional probability spaces via conditional UCWs. Next, we outline our conditional RIS (CRIS) that generalizes GRIS [[Lin et al. 2022](#)].

Take inputs $X_i \in \Omega_i$ with conditional UCWs $W_{X_i|Z}$, where samples X_i and domains Ω_i may both conditionally depend on Z . We must shift X_i into integrand f 's domain Ω with conditioned shift mappings, evaluating $Y_i = T_i(X_i|Z)$, e.g., transforming a neighbor pixel's suffix path to start from a path prefix in the current pixel (see [Section 6](#)).

Then we resample a Y by picking from Y_i in proportion to resampling weights $w_i = m_i(Y_i)\hat{p}(Y_i)W_{X_i|Z}|T_i'$, similarly to GRIS. This gives a (now conditional) unbiased contribution weight: $W_{Y|Z} = \frac{1}{\hat{p}(Y)} \sum w_i$, where w_i and \hat{p} may (implicitly) depend on Z . Together, Y and $W_{Y|Z}$ integrate properly in Y 's conditional support:

$$\mathbb{E}[f(Y)W_{Y|Z}|Z] = \int_{\text{supp}(Y|Z)} f(y) dy. \quad (13)$$

If seeking an unbiased estimate for the integral, we must ensure $\text{supp}(Y|Z)$ covers $\text{supp}(f)$, e.g., using a canonical input sample. Then, regardless of Z , we have $\mathbb{E}[f(Y)W_{Y|Z}|Z] = \int_{\Omega} f(y) dy$.

We can also use Y to resample iteratively, chaining multiple RIS passes to implement ReSTIR. We need to know $\text{supp}(Y|Z)$ to properly evaluate other samples' MIS weights. A canonical sample ensures $\text{supp}(Y|Z) = \text{supp}(\hat{p})$, with \hat{p} implicitly depending on Z .

5.2 Integration in the General Case

We now generalize Equation 12, borrowing samples from other domains by shift mapping. Our MIS weights use the different input domains' target functions \hat{p}_i (e.g., suffix path radiances) as proxies for unknown conditional PDFs.

We derive the formula by a shortcut: GRIS provides estimate $f(X)W_{X|Z}$ by selecting X randomly from the inputs, but we sum probability times contribution over the choices. Similarly to Lin et al.'s [2022] offline estimator, this reduces color noise, but is equivalent for scalar-valued f when choosing $\hat{p} = f$.

We assume M inputs X_i in different domains, conditioned by Z , and shift them into f 's domain via $Y_i = T_i(X_i|Z)$. The generalized conditional UCW estimator, given Z , is then

$$\langle I_Z \rangle = \sum_{i=1}^M m_i(Y_i|Z) f(Y_i|Z) W_{X_i|Z} |T'_i(X_i|Z)|, \quad (14)$$

for MIS weights m_i , integrand f (e.g., suffix radiance), pre-shift conditional UCW $W_{X_i|Z}$, and shift Jacobian determinant $|T'_i|$. Any MIS weights from Lin et al. [2022] can be used if we interpret the formulas with an implicit conditioning by Z , which we do from now on. We specifically mention the generalized balance heuristic,

$$m_i(y) = \frac{\alpha_i \hat{p}_{\leftarrow i}(y)}{\sum_{j=1}^M \alpha_j \hat{p}_{\leftarrow j}(y)}, \quad (15)$$

where α_i are domain weights specifying the relative weight (confidence) given to the samples, and

$$\hat{p}_{\leftarrow j}(y) = \hat{p}_j(T_j^{-1}(y)) |T_j'^{-1}(y)| \quad (16)$$

reads the (conditional) target function at y 's corresponding path in pixel j , i.e., $T_j^{-1}(y)$, as a proxy for its conditional PDF, and the Jacobian determinant transforms this proxy the same way probability densities transform in shift mappings.

This $\langle I_Z \rangle$ integrates over the samples' supports in f 's domain:

$$\mathbb{E}[\langle I_Z \rangle | Z] = \int_{\bigcup_{i=1}^M \text{supp}(Y_i|Z)} f(y) dy. \quad (17)$$

If the union covers the integrand, e.g., by including a canonical sample, this estimator is unbiased. We use this estimator to integrate suffix radiance by defining Z as the *supporting prefixes*, i.e., unused parts of our reused (sub)paths.

Next, we describe a framework using conditional RIS and ReSTIR in path tracing to reuse suffix paths between pixels and frames.

Algorithm 1: Pseudocode of our prototype final gather.

```

1 function SuffixReSTIR()
2   parallel foreach pixel  $q \in \text{Image}$  :
3      $q' \leftarrow \text{TemporalReprojection}(q)$ 
4      $X^P \leftarrow \text{TraceNewPrefix}(q)$  //  $\Downarrow$  Temporal prefix reuse w. GRIS
5     Reservoirs[ $q$ ]. $X^P \leftarrow \text{GRIS}(X^P,$ 
6       prevReservoirs[ $q'$ ]. $X^P$ ) // UCWs omitted for conciseness.
7     // Suffix reuse with conditional RIS (CRIS). Pass in prefixes too.
8      $X^S \leftarrow \text{TraceNewSuffix}(\text{Reservoirs}[q].X^P)$ 
9     Reservoirs[ $q$ ]. $X^S \leftarrow \text{CRIS}(X^S, \text{prevReservoirs}[q'].X^S)$ 
10    Reservoirs[ $q$ ]. $X^S \leftarrow \text{SpatialSuffixReuse}(\text{Reservoirs})$ 
11    prevReservoirs  $\leftarrow$  Reservoirs // Save for the next frame.
12    // Final gather phase, implements Equation 23.  $\Downarrow$  First prefix.
13     $(X^P, X^S) \leftarrow \text{TraceFullPath}(q)$ 
14    // Search for  $k = M - 1$  reservoirs with nearest supporting prefixes.
15     $[R_1, \dots, R_k] \leftarrow \text{FindSpatialKNN}(\text{Reservoirs}, X^P, k)$ 
16    Color[ $q$ ] += ComputeMIS( $[X^P, R_1.X^P, \dots, R_k.X^P], X^S$ )
17    * PathContrib( $X^P, X^S$ ) + Gather( $[R_1, \dots, R_k], X^P$ ) /  $N$ 
18    // Gather: MIS over  $X^P$  and all  $R_j.X^P$  and contribute all  $R_j.X^S$ 
19  for  $i \in [2, \dots, N]$  : //  $\Downarrow$  Other prefixes.
20     $X^P \leftarrow \text{TraceNewPrefix}(q)$ 
21     $[R_1, \dots, R_k] \leftarrow \text{FindSpatialKNN}(\text{Reservoirs}, X^P)$ 
22    Color[ $q$ ] += Gather( $[R_1, \dots, R_k], X^P$ ) /  $N$ 

```


6 SUFFIX RESTIR

Applying conditional RIS theory to unidirectional path tracing, we build a proof-of-concept prototype¹ that produces well-distributed *suffix* subpaths and reuses them spatiotemporally. We sample new prefixes at integration, connecting each such prefix to one or more conditionally-sampled suffixes. We update suffixes via ReSTIR, improving their distribution temporally. We summarize our prototype in Algorithm 1 (our supplemental document has more details). Note that reservoir sizes double, versus ReSTIR PT [Lin et al. 2022], as we must store both prefix and suffix data.

Reused suffixes may not entirely cover the suffix domains for new integration prefixes, so we combine with a canonical suffix sample via MIS, to guarantee coverage. This ensures unbiasedness.

Our figures show reconnection shifts with one-bounce prefixes, but we actually use Lin et al.'s [2022] hybrid shift, postponing reconnection on low-roughness vertices. Prefixes end at the second consecutive high-roughness vertex; the remainder is the suffix.

6.1 Prefix and Suffix Distributions



Reservoirs can reside in many domains, including world space, but we describe the concept in screen space, where each pixel has a *reservoir*, storing a (yellow) prefix and (red) suffix path. We call the prefix a *supporting prefix*, as it conditions the suffix random

¹Prototype code at: <https://github.com/NVlabs/conditional-restir-prototype>

variable, determining its support and path space coverage. The supporting prefix is needed for conditional MIS weights, shift mappings and target functions for reuse and integration.

Well-distributed supporting prefixes are vital if suffixes are to capture important light paths. This suggests regularly updating with new independent prefixes. But changing a prefix modifies its suffix’s support and target function, worsening its distribution. We strike a balance by updating prefixes temporally with ReSTIR, without spatial reuse. At each frame (Alg. 1, lines 2-5), we sample a new (blue) supporting prefix candidate for each pixel.

We find the matching pixel in the prior frame, and GRIS-select between this candidate and the old supporting prefix (yellow), shifted to the new pixel (dashed). We favor the old frame’s supporting prefix in GRIS to let the distribution improve over time.

Next (lines 6-7), we update suffixes to the current frame. With the supporting prefix (blue, left) already in the current frame, we take the (yellow) old prefix and its suffix (orange), and RIS between the shifted old suffix (orange, right) and a new canonical suffix (cyan). This uses conditional RIS between the reused suffix and new canonical suffix. We give old suffixes higher domain weight so the suffix distribution improves over time.

After updating suffixes temporally, we also reuse spatially (line 8). This is similar, but rather than using one projected pixel from last frame, we use conditional RIS to resample between the current suffix and numerous suffixes from the local pixel neighborhood.

6.2 Target Functions and MIS Weights

Resampling always requires target functions and MIS weights. To define these, we decompose the path contribution function $f(x)$

$$f(x^p, x^s) = f_p(x^p) f_{ps}(x^p, x^s) f_s(x^s), \quad (18)$$

where f_p is path throughput over prefix x^p , f_{ps} is the reconnection term including two BSDFs and visibility, and f_s is the throughput over suffix x^s (after reconnection) times the emitted radiance.

We resample supporting prefixes via ReSTIR with target function

$$\hat{p}(x^p) = f_p(x^p). \quad (19)$$

This is stable and gives better quality than including suffix data, e.g., $\hat{p}(x^p) = f(x^p, x^s)$; adding a suffix dependency also makes avoiding bias tricky, as suffixes end up conditioned by themselves, invalidating their unbiased contribution weights (see Section S.1.1). We use the generalized balance heuristic (Equation 15) to MIS between the new canonical prefix and prior frame prefix.

For spatiotemporal suffix reuse we define pixel i ’s target function as the path contribution after supporting prefix X_i^p :

$$\hat{p}_i(x^s) = f_{ps}(X_i^p, x^s) f_s(x^s). \quad (20)$$

Suffix x^s is conditioned on prefix X_i^p , so X_i^p is available as a constant. We leave f_p out, as all suffixes from X_i^p share this prefix and




Figure 2: Given per-pixel paths (left), ReSTIR PT reuses by modifying nearby paths to use the current primary hit (middle). We delay reuse (at least) to secondary hits (right), reducing correlation but increasing noise. We address noise by a final gather, i.e., averaging reuse at multiple short prefixes. Reused suffixes are conditioned on the removed random prefixes that determined their secondary hit locations.

$f_p(X_i^p)$ cancels out during resampling. For MIS weights, we again use the generalized balance heuristic (Equation 15), with all terms conditioned with the supporting prefixes. Dropping f_p also removes a source of imbalance in the MIS weights, reducing variance.

6.3 Integration with Borrowed Suffixes

At integration time, we sample an independent integration prefix X^p (blue) and search nearby for similar supporting prefixes (circled)

in world-space by examining path geometry and borrow their suffixes. We sample a (cyan) canonical suffix to combine its contribution with the borrowed suffixes. Our proof-of-concept uses distance between the end of our integration prefix and supporting prefixes to select suffixes for reuse, as overlapping domains are more likely; exploring other heuristics remains an open question.

To integrate, we substitute suffix contribution $f_{ps}(X^p, x^s) f_s(x^s)$ into Equation 14 (Z includes X^p and all supporting prefixes, which we keep implicit in the following). The estimator $\langle I_s \rangle$ for suffix contributions integrates over X^p ’s suffix space $\Omega^s(X^p)$:

$$\mathbb{E}[\langle I_s \rangle | X^p] = \int_{\Omega^s(X^p)} f_{ps}(X^p, x^s) f_s(x^s) dx^s. \quad (21)$$

We then multiply $\langle I_s \rangle$ by prefix throughput estimate $f_p(X^p) W_{X^p}$ to get the joint estimator for the full path integral:

$$\langle I \rangle = \sum_{i=1}^M m_i(Y_i^s) f(X^p, Y_i^s) W_{X^p, Y_i^s}, \quad (22)$$


where $Y_i^s = T_i(X_i^s)$ is the suffix X_i^s shifted to continue from prefix X^p , f is the full path contribution function, m_i weights the different suffixes, and $W_{X^p, Y_i^s} = W_{X_i^s | X^p} W_{X^p} |T_i'(X_i^s)|$ is the joint UCV (Section 4.1) for the full path after shifting the suffix, as the Jacobian transforms $W_{X_i^s | X^p} |T_i'(X_i^s)|$ into $W_{Y_i^s | X^p}$. One of the Y_i^s is the canonical suffix that guarantees full coverage of the suffix space.

7 FINAL GATHER

In Section 6 we described a way to distribute reusable path suffixes, integrating with short prefixes (Figure 2, right), finding supporting prefixes with similar last vertex geometry, and connecting to their resampled suffixes (red). Shrinking supporting prefix length to one simplifies to ReSTIR PT (middle). Distracting resampling artifacts

can arise if strong outliers get widely reused spatiotemporally. This impoverishes the sample pool, causing correlations. Combining a random prefix with reused subpaths increases path variety. But this adds noise, which we address with a final gather (Figure 2, right).


Photon mapping often uses a similar solution. Querying radiance at primary hits leads to blotchy artifacts. A *final gather* moves photon queries to the second diffuse hit.



Like radiance caching, suffix ReSTIR fills space with cache points (yellow), gathering by tracing integration prefixes (blue). Rather than shooting photons from lights, we sample supporting camera prefixes (yellow), concentrating them in high-throughput areas. Instead of interpolating radiances, we reuse resampled suffixes (red) by shift mapping them to the current integration prefix.

Estimates of pixel color are conditioned by the supporting prefixes chosen, and must individually be unbiased. As described in Section 6.3, we pick similar support-

ing prefixes to maximally match their covered path spaces with our integration prefix. Still, we cannot assume that borrowed suffixes alone cover the integration domain, so by default, each integration prefix must sample a canonical suffix. Suffix paths can be long, so this would be prohibitively expensive. We apply a form of Russian roulette to reap the benefits of ReSTIR suffixes without tracing more canonical suffixes.



Our final gather estimator builds on borrowed suffixes (Section 6.3) with the premise that integration prefixes are relatively cheap. We trace a number N of prefixes and average their estimators (Equation 22), each reusing the suffixes of its closest supporting prefixes. We conceptually include canonical suffixes for all N integration prefixes, specifically including them in our MIS weights—but, by Russian roulette, we replace canonical suffix contributions with zero except for one random prefix, whose contribution is multiplied by N . But, the multiplication cancels the mean, resulting in an unweighted contribution. As much of path space is likely covered by borrowed ReSTIR suffixes, this roulette greatly improves rendering efficiency. The canonical suffix is only needed for the non-covered minority.

Assuming symmetrically sampled prefixes X_i^p , we can choose prefix X_1^p to include canonical suffix X_{11}^s , with no shift mapping required. The other prefixes X_i^p only have reused suffixes $X_{i2}^s, \dots, X_{iM}^s$. This leads to the following final gather estimator:

$$\begin{aligned} \langle I_{FG} \rangle &= m_1(X_{11}^s | X_1^p) f(X_1^p, X_{11}^s) W_{X_1^p, X_{11}^s} \\ &+ \frac{1}{N} \sum_{i=1}^N \sum_{j=2}^M m_i(Y_{ij}^s | X_i^p) f(X_i^p, Y_{ij}^s) W_{X_i^p, Y_{ij}^s} \end{aligned} \quad (23)$$

with definitions analogous to Equation 22. The first line is the contribution from the canonical suffix X_{11}^s (lines 10-12 in Algorithm 1). The inner sum on the second line does the gather (lines 12-16).

ReSTIR-guided suffixes tend to be high quality, and the integration over the dimension freed by postponing reuse tends to be the

variance bottleneck. We seek ReSTIR PT’s low noise without its correlation and present a final gather as a candidate solution, but acknowledge our proof-of-concept approach leaves much yet to do.

8 PROOF-OF-CONCEPT EXPERIMENTS

To study our theory’s potential to improve quality without tracing many more independent paths, we built a proof-of-concept algorithm using conditional RIS and ReSTIR, per Sections 6 and 7. We started from Lin et al.’s [2022] ReSTIR PT code and Falcor’s path tracer [Kallweit et al. 2022]. We use RTXDI [NVIDIA 2021] for direct lighting and apply NEE at prefix vertices to account for lighting not covered by suffixes. Results use an NVIDIA RTX 4090 at 1920×1080 with max path length of 12. Our supplemental material contains more implementation and performance details, plus full images and a video studying temporal behavior.

We deem our work mostly theoretical. Experiments focus on uncovering insights on the benefits of conditional RIS theory for subpath reuse. Thus, we have not yet searched for optimally-performing algorithms or implementations. We show some equal-time tests, but our goal is not showing our prototype somehow faster or better, but understanding challenges and identifying promising future work.

All in all, our experiments suggest relatively high potential, with some high-reward research directions such as improving importance sampling of integration prefixes, e.g., with ReSTIR or path guiding.

Driving suffixes with ReSTIR helps greatly. In Figure 3 we compare our suffix reuse (“Ours”) to using independent suffixes (“MMIS”), with a form of marginal MIS [West et al. 2022]. We show an inset from VEACH AJAR; all light comes indirectly through a barely-open door. Spatiotemporal suffix reuse greatly improves image quality.

One reused suffix may be enough. Figure 3 also ablates over use of multiple path suffixes during integration. For MMIS, adding more suffixes greatly improves quality; our suffix reservoirs already aggregate multiple suffixes, so adding more provides diminishing returns (similar to Wyman and Pantelev [2021]).

Final gather is important. We postpone reuse by one path vertex (vs. ReSTIR PT) and use Monte Carlo integration on the freed dimensions. This increases noise; our prototype lowers this with a final gather. Figure 4 compares various integration prefix counts. Despite VEACH AJAR’s indirect lighting, increasing *path prefixes* may suffice! In some sense, we convert path tracing into one-bounce integration, as if using an unbiased radiance cache—except we reuse full paths. Integration prefix count gives a natural quality slider for our prototype, but in the future, we hope to approach the quality of many prefixes while using fewer, via better importance sampling.

Russian roulette in the final gather improves efficiency. Figure 4 also compares canonical suffix count. We test one canonical suffix per prefix (right column) versus just one per pixel via Russian roulette (other columns); roulette reduces ray count by 80% with often limited quality impact. Sometimes tracing more canonical suffixes could be worthwhile, especially for offline rendering. See the supplemental document for a study of convergence behavior.

Suffix ReSTIR helps with disocclusion and movement. Figure 5 compares our prototype to ReSTIR PT with camera motion in VEACH AJAR. Screen space ReSTIR suffers “variance lag” around disocclusions as reservoirs get reset where temporal reuse fails. Our prototype does not suffer this lag. Disocclusions still invalidate reservoirs, but when integrating, we find valid reuse candidates from farther away by world-space matching. In some sense, we move ReSTIR to object space but still concentrate computation in visible pixels.

We largely fix ReSTIR spatiotemporal correlations. Figure 1 explores correlations in the TOWER BRIDGE. We render a static video, copying horizontal slices of consecutive frames to rows in the (x, t) plots. This shows temporal and spatial correlations as vertical and horizontal blobs, respectively. ReSTIR PT’s strong spatiotemporal correlations [Sawhney et al. 2022] are not present in our result.

Removing correlations helps with denoising. In Figure 6 we apply a pre-release version of DLSS-RR [NVIDIA 2023] to both ReSTIR PT and our prototype. Despite more raw noise at equal time, the lack of correlations often allows the denoiser to produce better quality with our subpath reuse.

High potential with improved final gather. In Figure 6 we compare our current prototype to ReSTIR PT at equal-time (8 prefixes) and also with higher prefix counts. Figure 7 further studies future potential, assuming a high-quality final gather. ReSTIR suffixes often contain enough information to produce clearer images than ReSTIR PT, showing less coloration, noise, and edge artifacts. This highlights the importance of making final gather cheaper.

9 FUTURE WORK

We believe our new conditional RIS theory offers a promising path to extend real-time rendering to harder light paths. Below we list some interesting future research enabled by our theory:

More efficient final gather. Postponing reuse via a final gather largely eliminates ReSTIR correlations. But we need an *efficient* final gather for cheap, high-quality prefixes. Better importance sampling may be key: path guiding, more resampling steps, low-discrepancy samples, or better stratification might all offer improvements.

Choosing the right supporting prefixes. For quick prototyping, we used a BVH range search [Evangelou et al. 2021] to find supporting prefixes near our integration prefix. Faster ways to find candidates likely exist; better selection heuristics might also improve quality.

Other theoretical applications. Our prototype shows potential theoretical benefits in the context of a simple final gather. Conditional RIS may also have other uses, e.g., spectral rendering [Weidlich et al. 2022] or volume rendering. Improving specific algorithmic details may yield more practical variants, e.g., better shift maps, variable prefix lengths, multi-vertex reuse, mid-prefix direct lighting, etc.

10 CONCLUSION

We present a new conditional RIS theory, generalizing unbiased contribution weights to allow Monte Carlo integration and resampling even with unknown conditional and joint PDFs. This works for chained RIS, enabling resampling with conditional MIS weights and shift mappings, extending Lin et al. [2022].

We apply our theory to spatiotemporally aggregate suffix paths, driving the conditioning prefixes and reused suffixes by ReSTIR. This allows bidirectional-like path reuse with unidirectional paths, focusing computation in visible regions.

Our proof-of-concept, unbiased final gather combines cheap path prefixes with suffixes reused via *conditional RIS*. This fills path space with unbiased cached suffixes analogous to photons, but sampled from the camera. Debiasing normally requires many canonical paths; Russian roulette allows skipping all but one. All together, this turns the renderer into a low-dimensional integration over short prefixes. But our prototype remains expensive; we need further algorithmic development for practical applications of conditional resampling.

ACKNOWLEDGMENTS

We want to thank Aaron Lefohn for discussions on and support of this research. Additional thanks to Matt Pharr for discussions and feedback on early ideas and paper drafts. Also, thank you to the anonymous reviewers for pointing out areas for clarification and improvement.

REFERENCES

- Pablo Bauszat, Victor Petitjean, and Elmar Eisemann. 2017. Gradient-Domain Path Reusing. *ACM Transactions on Graphics* 36, 6 (2017), 229:1–229:9. <https://doi.org/10.1145/3130800.3130886>
- Philippe Bekaert, Mateu Sbert, and John H Halton. 2002. Accelerating Path Tracing by Re-Using Paths. In *Eurographics Workshop on Rendering*. 125–134. <https://doi.org/10.2312/EGWR.EGWR02.125-134>
- Nikolaus Binder, Sascha Fricke, and Alexander Keller. 2019. Massively parallel path space filtering. (2019). arXiv:1902.05942 [cs.GR]
- 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* 39, 4 (2020), 148:1–148:17. <https://doi.org/10.1145/3386569.3392481>
- Min-Te Chao. 1982. A general purpose unequal probability sampling plan. *Biometrika* 69, 3 (1982), 653–656. <https://doi.org/10.2307/2336002>
- Xi Deng, Miloš Hašan, Nathan Carr, Zexiang Xu, and Steve Marschner. 2021. Path Graphs: Iterative Path Space Filtering. *ACM Transactions on Graphics* 40, 6 (2021), 276:1–276:15. <https://doi.org/10.1145/3478513.3480547>
- Xi Deng, Shaojie Jiao, Benedikt Bitterli, and Wojciech Jarosz. 2019. Photon Surfaces for Robust, Unbiased Volumetric Density Estimation. *ACM Transactions on Graphics* 38, 4 (2019), 46:1–46:12. <https://doi.org/10.1145/3306346.3323041>
- Luc Devroye. 1986. *Non-Uniform Random Variate Generation*. Springer-Verlag. <https://doi.org/10.1007/978-1-4613-8643-8>
- Victor Elvira, Luca Martino, David Luengo, and Mónica F. Bugallo. 2019. Generalized Multiple Importance Sampling. *Statistical Science* 34, 1 (2019), 129–155. <https://doi.org/10.1214/18-STS668>
- Jordanis Evangelou, Georgios Papaioannou, Konstantinos Vardis, and Andreas A. Vasilakis. 2021. Fast radius search exploiting ray-tracing frameworks. *Journal of Computer Graphics Techniques* 10, 1 (2021). <https://jcgf.org/published/0010/01/02/>
- Binh-Son Hua, Adrien Gruson, Victor Petitjean, Matthias Zwicker, Derek Nowrouzezahrai, Elmar Eisemann, and Toshiya Hachisuka. 2019. A Survey on Gradient-Domain Rendering. 38, 2 (2019), 455–472. <https://doi.org/10.1111/cgf.13652>
- Henrik Wann Jensen. 2001. *Realistic Image Synthesis Using Photon Mapping*. A. K. Peters, Ltd., USA.
- Simon Kallweit, Petrik Clarberg, Craig Kolb, Tomáš Davidovič, Kai-Hwa Yao, Theresa Foley, Yong He, Lifan Wu, Lucy Chen, Tomas Akenine-Möller, Chris Wyman, Cyril Crassin, and Nir Benty. 2022. The Falcor Rendering Framework. <https://github.com/NVIDIAGameWorks/Falcor> [Online; accessed 22-August-2023].
- Alexander Keller, Ken Dahm, and Nikolaus Binder. 2014. Path Space Filtering. In *ACM SIGGRAPH Talks*. Article 68. <https://doi.org/10.1145/2614106.2614149>
- Eric P. Lafortune and Yves D. Willems. 1993. Bi-Directional Path Tracing. In *Proceedings of Compugraphics*. 145–153. <https://graphics.cs.kuleuven.be/publications/BDPT/>
- Daqi Lin, Markus Kettunen, Benedikt Bitterli, Jacopo Pantaleoni, Cem Yuksel, and Chris Wyman. 2022. Generalized Resampled Importance Sampling: Foundations of ReSTIR. *ACM Transactions on Graphics* 41, 4 (2022), 75:1–75:23. <https://doi.org/10.1145/3528223.3530158>
- Daqi Lin, Chris Wyman, and Cem Yuksel. 2021. Fast Volume Rendering with Spatiotemporal Reservoir Resampling. *ACM Transactions on Graphics* 40, 6 (2021),

- 279:1–279:16. <https://doi.org/10.1145/3478513.3480499>
- Zander Majercik, Thomas Müller, Alexander Keller, Derek Nowrouzezahrai, and Morgan McGuire. 2021. Dynamic Diffuse Global Illumination Resampling. In *ACM SIGGRAPH Talks*. Article 24. <https://doi.org/10.1145/3450623.3464635>
- NVIDIA. 2021. RTX Direct Illumination API. <https://developer.nvidia.com/rtxdi>. [Online; accessed 22-August-2023].
- NVIDIA. 2023. NVIDIA DLSS 3.5: Enhancing Ray Tracing With AI. <https://www.nvidia.com/en-us/geforce/news/nvidia-dlss-3-5-ray-reconstruction>. [Online; accessed 12-September-2023].
- Yaobin Ouyang, Shiqiu Liu, Markus Kettunen, Matt Pharr, and Jacopo Pantaleoni. 2021. ReSTIR GI: Path Resampling for Real-Time Path Tracing. *Computer Graphics Forum* 40, 8 (2021), 17–29. <https://doi.org/10.1111/cgf.14378>
- Orion Pobursky. 2021. Tower Bridge Model (CC BY-2.0) available on LDraw.org. <http://omr.ldraw.org/files/1692> [Online; accessed 22-August-2023].
- Rohan Sawhney, Daqi Lin, Markus Kettunen, Benedikt Bitterli, Ravi Ramamoorthi, Chris Wyman, and Matt Pharr. 2022. Decorrelating ReSTIR Samplers via MCMC Mutations. arXiv:2211.00166 [cs.GR]
- Benjamin Segovia, Jean Claude Iehl, Richard Mitanchey, and Bernard Péroche. 2006. Bidirectional Instant Radiosity. In *Eurographics Symposium on Rendering*. 389–397. <https://doi.org/10.2312/EGWR/EGSR06/389-397>
- Justin Talbot, David Cline, and Parris Egbert. 2005. Importance Resampling for Global Illumination. In *Eurographics Symposium on Rendering*. 139–146. <https://doi.org/10.2312/EGWR/EGSR05/139-146>
- Lorenzo Tessari, Johannes Hanika, and Carsten Dachsbacher. 2017. Local quasi-monte carlo exploration. In *Eurographics Symposium on Rendering: Experimental Ideas & Implementations*. 71–81. <https://doi.org/10.2312/sre.20171196>
- Eric Veach and Leonidas Guibas. 1995a. Bidirectional estimators for light transport. In *Photorealistic Rendering Techniques*. Springer, 145–167. https://doi.org/10.1007/978-3-642-87825-1_11
- Eric Veach and Leonidas J. Guibas. 1995b. Optimally Combining Sampling Techniques for Monte Carlo Rendering. In *Proceedings of the 22nd Annual Conference on Computer Graphics and Interactive Techniques*. 419–428. <https://doi.org/10.1145/218380.218498>
- Eric Veach and Leonidas J. Guibas. 1997. Metropolis Light Transport. In *Proceedings of the 24th Annual Conference on Computer Graphics and Interactive Techniques*. 65–76. <https://doi.org/10.1145/258734.258775>
- Andrea Weidlich, Chloe LeGendre, Carlos Aliaga, Christophe Hery, Jean-Marie Aubry, Jiří Vorba, Daniele Siragusano, and Richard Kirk. 2022. Practical Aspects of Spectral Data in Digital Content Production. In *ACM SIGGRAPH Courses*. Article 11. <https://doi.org/10.1145/3532720.3535632>
- Rex West, Iliyan Georgiev, Adrien Gruson, and Toshiya Hachisuka. 2020. Continuous multiple importance sampling. *ACM Transactions on Graphics* 39, 4 (2020), 136:1–136:12. <https://doi.org/10.1145/3386569.3392436>
- Rex West, Iliyan Georgiev, and Toshiya Hachisuka. 2022. Marginal Multiple Importance Sampling. In *ACM SIGGRAPH Asia Conference Papers*. 42:1–42:8. <https://doi.org/10.1145/3550469.3555388>
- Chris Wyman and Alexey Pantelev. 2021. Rearchitecting Spatiotemporal Resampling for Production. In *ACM/EG Symposium on High Performance Graphics*. 23–41. <https://doi.org/10.2312/hpg.20211281>




Figure 3: We compare our ReSTIR-driven suffixes (e–h) to a variant without CRIS, tracing new suffixes every frame (a–d, “MMIS”). All results use four integration prefixes. We vary the suffix count reused for each integration prefix in the final gather. Without ReSTIR-driven suffixes we need many more for good quality (a–d), and the ray count increases quadratically due to the balance heuristic. ReSTIR suffixes are better distributed and give good results even with one suffix (e), avoiding the quadratic cost; Figure 4 shows that increasing prefix count is more cost-effective.

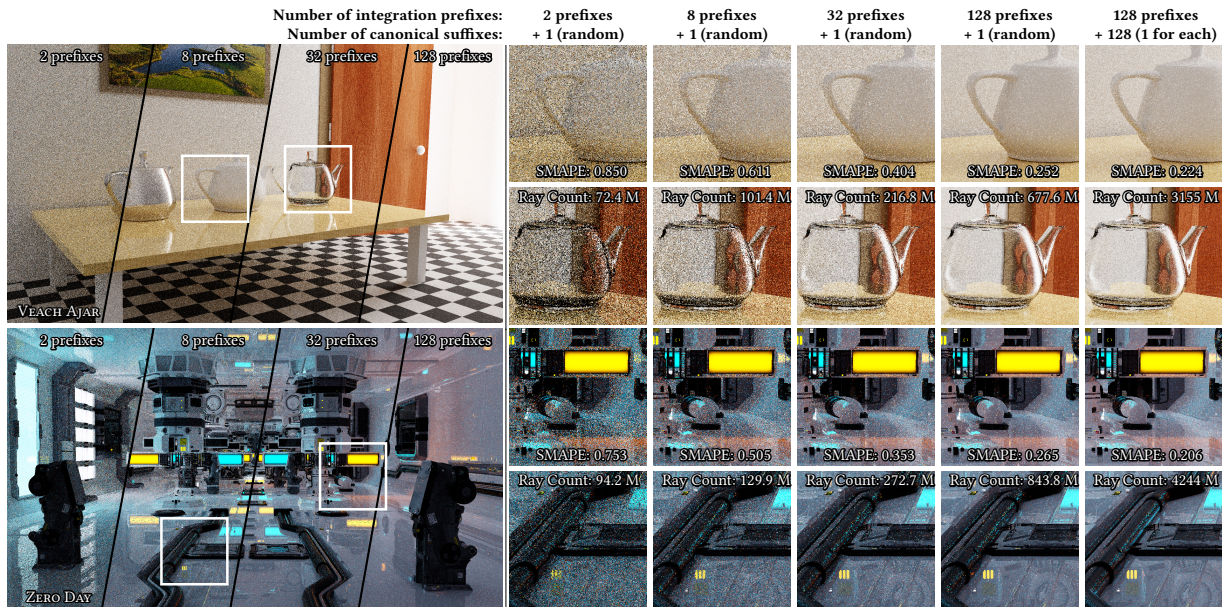


Figure 4: Effect of increasing integration prefixes. Here, each final gather prefix connects to one ReSTIR-driven suffix found via nearest-neighbor search. Image quality improves steadily with increased integration prefix count. We produce a canonical suffix for only one prefix, using Russian roulette, reducing ray count up to 80%. Quality loss from roulette is typically minor (compare the right two columns), except on some glossy surfaces.




Figure 5: Comparing disocclusion artifacts in our prototype and ReSTIR PT using fast camera motion. Because our final gather searches nearby candidates for suffix reuse based on vertices later in the path (not the primary hit point), it avoids ReSTIR PT’s discarded history near screen-space disocclusions. Both methods take one full path sample per pixel per frame, but feature additional rays for reuse.

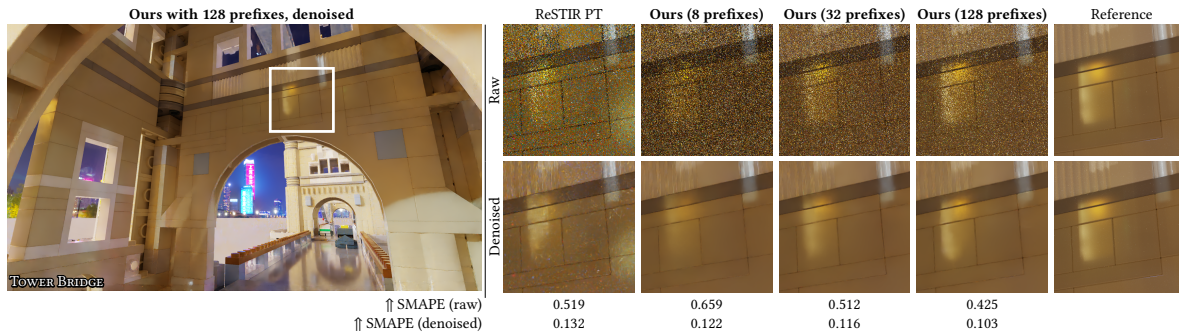


Figure 6: Our final gather produces better denoised results than ReSTIR PT, thanks to reduced correlation. Our unoptimized prototype with 8 integration prefixes (74 ms) is about equal-time to ReSTIR PT (76 ms) with increased candidate samples, giving a lower bound on achievable quality at interactive frame rates. Our prototype with 32 prefixes (207 ms) and 128 prefixes (736 ms) achieves superior quality but are more expensive. Future research should increase the achievable prefix count, with potential importance sampling improvements giving a multiplicative effect on the effective count.



Figure 7: Comparing path tracing, MMIS, ReSTIR PT, and our proof-of-concept with one full path per frame, assuming high-quality final gather. Our prototype and MMIS are here configured roughly equal time, which is much longer than ReSTIR PT’s cost; this figure estimates the potential opened by subpath reuse like a final gather. See Figure 6 for an equal-time comparison. We use 128 prefix samples connecting to one ReSTIR suffix per prefix; MMIS uses 64 prefixes connecting to three suffixes for an improved balance. With subpath reuse, we avoid color shift or correlation artifacts common in ReSTIR PT. All scenes feature a moderately-fast animated camera to prevent over-relying on temporal accumulation. See the supplemental document for the timings.