Marginal Multiple Importance Sampling

REX WEST, The University of Tokyo, Japan
ILIYAN GEORGIEV, Autodesk, United Kingdom
TOSHIYA HACHISUKA, University of Waterloo, Canada

Multiple importance sampling (MIS) is a powerful tool to combine different sampling techniques in a provably good manner. MIS requires that the techniques' probability density functions (PDFs) are readily evaluable point-wise. However, this requirement may not be satisfied when (some of) those PDFs are marginals, i.e., integrals of other PDFs. We generalize MIS to combine samples from such marginal PDFs. The key idea is to consider each marginalization domain as a continuous space of sampling techniques with readily evaluable (conditional) PDFs. We stochastically select techniques from these spaces and combine the samples drawn from them into an unbiased estimator. Prior work has dealt with the special cases of multiple classical techniques or a single marginal one. Our formulation can handle mixtures of those.

We apply our *marginal MIS* formulation to light-transport simulation to demonstrate its utility. We devise a *marginal path sampling* framework that makes previously intractable sampling techniques practical and significantly broadens the path-sampling choices beyond what is presently possible. We highlight results from two algorithms based on marginal MIS: a novel formulation of path-space filtering at multiple vertices along a camera path and a similar filtering method for photon-density estimation.

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

Additional Key Words and Phrases: rendering, multiple importance sampling, path filtering

ACM Reference Format:
Rex West, Iliyan Georgiev, and Toshiya Hachisuka. 2022. Marginal Multiple Importance Sampling. In *SIGGRAPH Asia 2022 Conference Papers (SA '22 Conference Papers), December 6–9, 2022, Daegu, Republic of Korea.* ACM, New York, NY, USA, 7 pages. https://doi.org/10.1145/3550469.3555388

1 INTRODUCTION

Continuous multiple importance sampling (CMIS) [West et al. 2020] generalizes multiple importance sampling (MIS) [Veach and Guibas 1995] to support uncountably infinite sets of sampling techniques. West et al. [2020] showed that CMIS is equivalent to sampling from the marginal distribution over the space of sampling techniques (*technique space*), and thereby the continuous generalization of the balance heuristic must evaluate the marginal PDF. This marginal PDF is often computationally intractable, so they proposed to stochastically sample a finite number of techniques from the technique space and make an approximation. This stochastic approximation, referred to as stochastic MIS (SMIS), yields an unbiased estimator at the cost of additional variance. However, SMIS, while powerful, is limited to combining samples drawn from techniques that are defined (parameterized) in a *single* technique space.

We introduce *marginal MIS* (MMIS), a generalization of SMIS to *multiple* technique spaces. Our key insight is that the marginal distribution over a technique space can be thought of as a single sampling technique with a marginal PDF (*marginal technique*), and multiple marginal techniques can be combined in a manner akin to the classical multi-sample MIS estimator. The result is a practical unbiased estimator combining multiple computationally intractable marginal distributions.

We demonstrate the practicality of MMIS with an application to path sampling in light-transport simulation. Path sampling is usually done by sequentially generating the vertices along a path, starting either from the sensor, a light source, or both. To define an estimator for an image pixel, we generally need to be able to evaluate the PDF for each sampled path. This requirement of having computationally tractable PDFs has severely limited the types of path sampling techniques we can use in practice. We devise a *marginal path sampling* framework, based on MMIS, which enables the use of path sampling techniques with computationally intractable marginal PDFs. For example, such PDFs can arise from performing multiple connections to multiple subpaths to form a complete path, something that is beyond the capability of existing (bidirectional) sampling methods.

We present concrete examples that demonstrate how our path sampling framework allows formulating techniques that were previously not possible. We first augment the path-filtering method of West et al. [2020] to filter at multiple path vertices instead of only one. Unlike prior work [Deng et al. 2021] which builds on the concept of path graphs and fixed-point iteration to achieve multi-vertex filtering, our framework allows us to directly define a Monte Carlo estimator out of marginal path sampling techniques. We then show how our multi-vertex filtering formulation can be trivially applied to photon-density estimation, extending the benefits to scenes that require sampling from emitters. These new methods are possible only with our new framework which enables us to use samples from marginal distributions.

In summary, our contributions are as follows:

- marginal MIS framework that generalizes SMIS to combine samples drawn from multiple different technique spaces;
- path sampling framework that leverages this estimator to enable the use of a large arsenal of previously intractable path sampling techniques;
- path-integral formulation of multi-vertex path filtering; and
- multi-vertex filtering extension to photon-density estimation.

SA '22 Conference Papers, December 6–9, 2022, Daegu, Republic of Korea
© 2022 Copyright held by the owner/author(s). Publication rights licensed to ACM. This is the author's version of the work. It is posted here for your personal use. Not for redistribution. The definitive Version of Record was published in *SIGGRAPH Asia 2022 Conference Papers (SA '22 Conference Papers), December 6–9, 2022, Daegu, Republic of Korea*, https://doi.org/10.1145/3550469.3555388.

2 RELATED WORK

Veach and Guibas [1995] proposed multiple importance sampling (MIS) to combine a finite set of estimators with different PDFs. They refer to each PDF and its sampling procedure as a *sampling technique*. West et al. [2020] generalized MIS to handle an uncountably infinite set of sampling techniques, referred to as a *technique space*. They also proposed the practical stochastic MIS (SMIS) estimator which randomly selects a finite number of techniques from this continuous space. The key difference in our work is that we now consider multiple technique spaces in our estimator, instead of the single technique space in SMIS.

We focus on applications of our proposed framework to path reuse and filtering. Bekaert et al. [2002] introduced the first path reuse method that amortizes path sampling cost by deterministically tracing a ray for every pixel in a tile, and connecting each to a single, shared suffix path. Bauszat et al. [2017] leveraged recent advances in gradient-domain rendering and deterministic shift mapping to support a wider range of materials and improve the convergence rate of of path reuse by an order of magnitude. Our marginal path sampling framework generalizes the concept of path reuse beyond those cases where the resulting path probability density is readily computable.

Path-space filtering [Keller et al. 2014] refines the contributions of sampled paths by connecting a prefix of one path to suffixes of other paths whose prefixes lie in close proximity. West et al. [2020] showed that this process can be formulated as multiple importance sampling over a continuous technique space where prefixes identify techniques and suffixes represent samples from those techniques, enabling unbiased path filtering for the first time. Deng et al. [2021] extended path filtering to multiple vertices along a path by reformulating filtering as the fixed point solution to a linear system of aggregation and propagation. While their numerical results support the idea of the fixed-point iteration, the link between this iterative process and the underlying transport equations is not fully clear. We show how our marginal path sampling can be applied to the same problem to derive a path-based formulation for multi-vertex path filtering. Our formulation is based on path sampling from marginal distributions and is an explicit estimator of a path integral. We further show the relationship between our formulation and the iterative process proposed by Deng et al. [2021] in Section 5.

3 MARGINAL MULTIPLE IMPORTANCE SAMPLING

We begin by deriving a *marginal MIS* (MMIS) estimator that combines sampling techniques with marginal densities. Using such a marginal density requires typically intractable marginalization over a space of continuous, auxiliary variables that condition sampling. In contrast to West et al. [2020]'s SMIS, our MMIS considers multiple such spaces of conditioning variables, i.e. technique spaces, one for each marginal density, with potentially differing dimensionality.

Consider the integral $I = \int_{\mathcal{X}} f(x)\,\mathrm{d}x$ of a real-valued function f over a domain $\mathcal{X}$ and a multi-sample MIS estimator for it that draws n_i samples from each of $T \geq 1$ sampling techniques with probability density function (PDFs) p_i and associated weighting function w_i:

$$\langle I \rangle_{\mathrm{MIS}} = \sum_{i=1}^{T} \sum_{j=1}^{n_i} \frac{w_i(x_{i,j}) f(x_{i,j})}{n_i p_i(x_{i,j})} = \sum_{i=1}^{T} \sum_{j=1}^{n_i} \langle I \rangle_{\mathrm{MIS}}^{i,j}. \tag{1}$$

This estimator can be directly used when each PDF p_i can be evaluated exactly. However, $p_i(x)$ could be a *marginal* density,

$$p_i(x) = \int_{\mathcal{T}_i} p_i(x,t)\mathrm{d}t = \int_{\mathcal{T}_i} p_i(x|t) p_i(t)\mathrm{d}t, \tag{2}$$

which is the expectation of a conditional density $p_i(x|t)$ over some auxiliary variable $t \in \mathcal{T}_i$ with density $p_i(t)$. Such a marginal PDF arises when, for example, a sampling procedure first chooses a random t value which it then uses to condition the sampling of x (e.g., choosing the standard deviation t of a Gaussian distribution and then sampling x from it). We will refer to techniques with readily computable PDFs $p_i(x)$ as *classical*, and techniques requiring marginalization over auxiliary variables we will call *marginal*.

The marginal integral (2) generally has no closed-form solution. This hampers the practical use of marginal techniques, since even unbiased estimation of the integral, which appears in the denominator of Eq. (1), generally leads to a biased MIS estimator [West et al. 2020]. We will address this problem via an approximation of the MIS estimator that yields sub-optimal yet *unbiased* estimates of the sought integral.

3.1 Marginal MIS estimator

When using the balance heuristic [Veach and Guibas 1995], the MIS weighting functions w_i for marginal techniques take the form

$$w_i(x) = \frac{n_i p_i(x)}{\sum_{i'=1}^{T} n_{i'} p_{i'}(x)} = \frac{n_i p_i(x)}{\sum_{i'=1}^{T} n_{i'} \int_{\mathcal{T}_{i'}} p_{i'}(x|t) p_{i'}(t)\,\mathrm{d}t}. \tag{3}$$

Plugging w_i into the MIS estimator (1) and simplifying it yields

$$\langle I \rangle_{\mathrm{MIS}}^{i,j} = \frac{f(x_{i,j})}{\sum_{i'=1}^{T} n_{i'} \int_{\mathcal{T}_{i'}} p_{i'}(x_{i,j}|t) p_{i'}(t)\,\mathrm{d}t}. \tag{4}$$

Next, recall that sampling from a marginal technique involves drawing an auxiliary variable $t \sim p_i(t)$ which conditions the sampling of $x \sim p_i(x|t)$. For each marginal technique i, we generate n_i such pairs $(t_{i,j}, x_{i,j})$, with $j = 1..n_i$. Following West et al. [2020], we reuse those pairs to construct an n_i-sample estimate of the technique's marginal integral in Eq. (4):

$$\langle I \rangle_{\mathrm{MIS}}^{i,j} \approx \frac{f(x_{i,j})}{\sum_{i'=1}^{T} \cancel{n_{i'}} \left[\frac{1}{\cancel{n_{i'}}} \sum_{j'=1}^{n_{i'}} \frac{p_{i'}(x_{i,j'}|t_{i',j'}) \cancel{p_{i'}(t_{i',j'})}}{\cancel{p_{i'}(t_{i',j'})}} \right]}. \tag{5}$$

Plugging this result back into Eq. (1), we obtain our *marginal MIS* (MMIS) estimator:

$$\langle I \rangle_{\mathrm{MMIS}} = \sum_{i=1}^{T} \sum_{j=1}^{n_i} \frac{f(x_{i,j})}{\sum_{i'=1}^{T} \sum_{j'=1}^{n_{i'}} p_{i'}(x_{i,j}|t_{i',j'})}. \tag{6}$$

This estimator is an *approximation* of Eq. (1), yet it is an *unbiased* estimator for the sought integral I, as long as at least one sampled conditional technique has positive density for every point x where $f(x) > 0$.

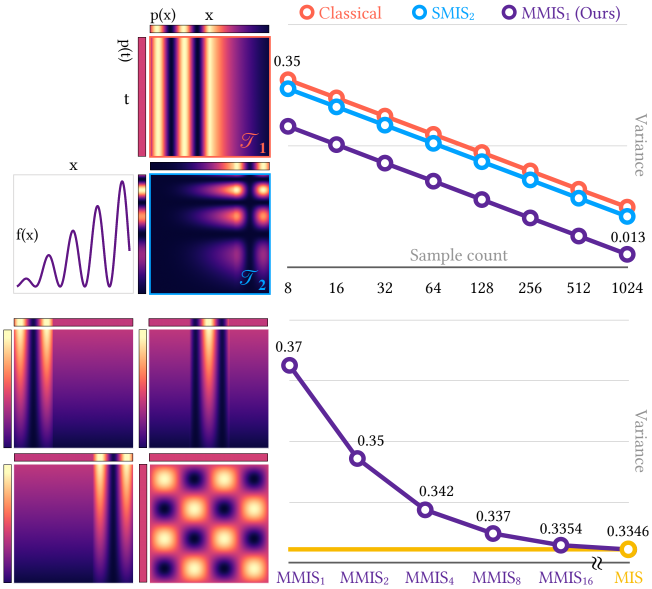

Fig. 1. Top: Our MMIS estimator can mix classical (top square) and marginal (bottom square) techniques to achieve lower variance than prior theory can. Bottom: MMIS approximates a (hypothetical) MIS estimator that evaluates marginal PDFs (here, four) exactly. The additional variance vanishes the number of techniques sampled from each technique space increases.

Note that we do not need to use all available (marginal) techniques in the estimation. We are free to choose a (random) subset, and so long as the above unbiasedness conditions hold, the MMIS remains unbiased. Note also that for $T = 1$ the MMIS estimator reduces to the SMIS estimator of West et al. [2020].

The estimator in Eq. (6) involves only the conditional sampling PDFs $p_i(x|t)$ which must be readily computable—a reasonable requirement in path sampling as we will discuss later. Similarly to West et al. [2020], we think of the auxiliary variable t as a conditional-technique identifier that lives in an associated *technique space* $\mathcal{T}_i$.

Combining classical and marginal techniques. For simplicity, we derive the MMIS estimator (6) from a variant of the balance heuristic (3) that assumes all techniques are marginal. However, mixing in classical techniques (i.e., with readily computable densities) is straightforward. Each classical technique can be extended to an arbitrary space of conditional techniques that are all identical; Fig. 1 (top) shows an example. Alternatively, one can split the PDF sum in Eq. (3) into a sum of classical and a sum of marginal PDFs. Following through the derivation, only the PDFs in the latter sum are expanded into marginal integrals that require MMIS.

3.2 Canonical experiments

In Fig. 1 we perform two experiments that illustrate the properties and trade-offs of our MMIS estimator (6). We plot the variance of estimating a 1D function $f(x)$ using marginal techniques with a 1D technique space each. Every 2D plot shows a joint density $p(x, t)$ along with its 1D marginals $p(x)$ and $p(t)$.

With only the MIS and SMIS theories, we must decide whether to construct an estimator using a set of classical techniques or using a single marginal technique, respectively. In Fig. 1, top, we demonstrate how our marginal MIS allows mixing classical and marginal techniques. The top PDF represents a classical technique that importance samples the left part of the integrand well. The bottom is a marginal technique whose (intractable) PDF $p(x)$ importance samples the right part of the integrand. By extruding the classical technique over a technique space, and applying our marginal MIS, we combine the two into a single estimator that importance samples the entire integrand well and outperforms both sampling using only the classical technique and SMIS with only one marginal technique.

Figure 1, bottom, shows four marginal techniques. The first three achieve good importance sampling of different integrand features but are by individually biased, i.e., do not each cover the entire domain. The fourth technique has a defensive uniform distribution for x. We measure the variance of a reference (in normal circumstances infeasible) classical MIS estimator (1) that analytically evaluates the x-marginal of each technique and several variants of our MMIS approximation (6) of that estimator. Each estimator is given a total budget of $N = 64$ samples, where MMIS_n draws n technique-sample pairs (i.e., n_i in Eq. (6)) from each of the T marginal techniques and averages over $N/(nT)$ independent realizations. The additional variance incurred by our MMIS approximation vanishes as n increases at the cost of computational cost.For a fixed total sample count N the total number of PDF evaluations in MMIS_n is nNT, i.e., linear in n. In contrast, the number of PDF evaluations NT in the reference MIS estimator is independent of the sample allocation (and so is its variance). Determining the optimal choice of n in MMIS is problem-specific as it depends on the cost of PDF evaluation.

4 PATH SAMPLING USING MARGINAL MIS

Light-transport simulation using Monte Carlo integration involves sampling light paths connecting the emitters in a scene to sensors. In this section, we explain how our MMIS enables unbiased simulation in ways not possible with the existing path sampling frameworks. Specifically, it allows us to employ sampling techniques with previously intractable marginal PDFs.

4.1 Monte Carlo path sampling

The classical path-integral formulation of light transport expresses the value of an image pixel as a conceptually simple integral over the space $\mathcal{P}$ of all possible paths in the scene [Veach 1997]. In practice this integral is estimated using Monte Carlo sampling:

$$I = \int_{\mathcal{P}} f(\overline{\mathbf{x}}) \mathrm{d}\overline{\mathbf{x}} \approx \frac{1}{n} \sum_{i=1}^{n} \frac{f(\overline{\mathbf{x}}_i)}{p(\overline{\mathbf{x}}_i)}, \tag{7}$$

where $f(\overline{\mathbf{x}})$ is the energy contribution of path $\overline{\mathbf{x}} = \mathbf{x}_1 \ldots \mathbf{x}_k$ with k vertices. The sampling density $p(\overline{\mathbf{x}}) = p(\mathbf{x}_1, \ldots, \mathbf{x}_k)$ of a path is the joint density of its vertices. Each unique way to construct a given path constitutes a *path sampling technique* with a distinct PDF.

To estimate the pixel integral (7) we need to sample paths and evaluate their PDFs. The most commonly used sampling technique is path tracing [Kajiya 1986] which generates vertices sequentially, starting from the sensor (see figure below).

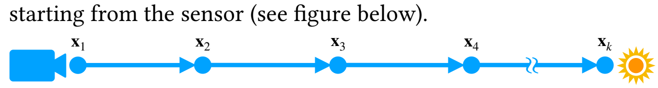

The PDF for sampling a k-vertex path in this manner is

$$p(\bar{\mathbf{x}}) = p(\mathbf{x}_1, \ldots, \mathbf{x}_k) = p(\mathbf{x}_1)p(\mathbf{x}_2|\mathbf{x}_1)\prod_{i=3}^{k} p(\mathbf{x}_i|\mathbf{x}_{i-1}, \mathbf{x}_{i-2}) \,. \tag{8}$$

Since each vertex is sampled dependent only on the preceding two vertices, and all conditioning variables are part of the output path sample, the resulting path PDF is easy to evaluate.

4.2 Problem statement

In theory, one has the full freedom to construct paths, i.e., vertex sequences, in arbitrary ways. Unbiased pixel estimation has traditionally been possible only if the technique's associated PDF is *readily computable*, since it is required by the pixel estimator (7) or an MIS variant thereof. This requirement has significantly limited the set of practically usable techniques.

To illustrate why this restriction is severe, consider a slight modification to the path-tracing technique where we "skip" the second sampled vertex, i.e., we exclude $\mathbf{x}_2$ from the constructed path $\bar{\mathbf{x}}$. Sampling proceeds as before, with vertices still conditioned on $\mathbf{x}_2$.

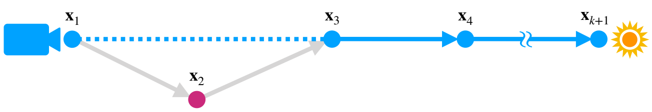

The vertex $\mathbf{x}_2$ will thus not show up in the PDF of the path as it is no longer a part of it:

$$p^*(\bar{\mathbf{x}}) = p^*(\mathbf{x}_1, \mathbf{x}_3, \ldots, \mathbf{x}_{k+1}) \tag{9}$$

$$= p^*(\mathbf{x}_1)p^*(\mathbf{x}_3|\mathbf{x}_1)p^*(\mathbf{x}_4|\mathbf{x}_3)\prod_{i=5}^{k+1} p^*(\mathbf{x}_i|\mathbf{x}_{i-1}, \mathbf{x}_{i-2}) \,. \tag{10}$$

The vertices $\mathbf{x}_3$ and $\mathbf{x}_4$ are sampled conditionally on $\mathbf{x}_2$, but their PDFs $p^*(\mathbf{x}_3|\mathbf{x}_1)$ and $p^*(\mathbf{x}_4|\mathbf{x}_3)$ are not conditioned on it as they marginalize it out:

$$p^*(\mathbf{x}_3|\mathbf{x}_1) = \int_{\mathcal{M}} p^*(\mathbf{x}_3|\mathbf{x}_2, \mathbf{x}_1)p^*(\mathbf{x}_2|\mathbf{x}_1)\mathrm{d}\mathbf{x}_2 \,, \tag{11}$$

$$p^*(\mathbf{x}_4|\mathbf{x}_3) = \int_{\mathcal{M}} p^*(\mathbf{x}_4|\mathbf{x}_3, \mathbf{x}_2)p^*(\mathbf{x}_2|\mathbf{x}_1)\mathrm{d}\mathbf{x}_2 \,, \tag{12}$$

where $\mathcal{M}$ is the union of all scene (surface) points. Recall that the computation of these PDFs is necessary to evaluate the pixel estimator (7). Simply estimating the marginal integrals would bias the estimator [Qin et al. 2015].

In general, marginals appear whenever a technique samples vertices conditionally on other vertices that are not included in the output path. Consequently, such techniques have not been used in light-transport simulation. Next we will show how our marginal MIS enables unbiased pixel estimation using such techniques.

4.3 A marginal path sampling framework

The PDF p^* in Eq. (9) belongs to a class of sampling techniques where the *sample vertices* $\bar{\mathbf{x}}$ are drawn conditionally on a set of *auxiliary vertices* $\bar{\mathbf{t}}$ that we need to marginalize over. For that technique we can visualize this split as in the inline figure below.

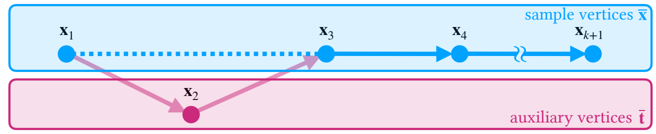

We will refer to the class of techniques that conform to this sampling model as *marginal path sampling techniques*. These techniques have a clear correspondence to the marginal techniques of marginal MIS (6), $p_i(x_{i,j}|t_{i,j}) \longrightarrow p_i(\bar{\mathbf{x}}_{i,j}|\bar{\mathbf{t}}_{i,j})$, where the sample variable x and technique variable t are replaced by their path-sampling counterparts $\bar{\mathbf{x}}$ and $\bar{\mathbf{t}}$. The technique space $\mathcal{T}$ then naturally arises as the space of possible values for the auxiliary vertices $\bar{\mathbf{t}}$. For the technique in Eq. (9), this space is $\mathcal{M}$. Classical techniques do not utilize auxiliary vertices, i.e., $\bar{\mathbf{t}} = \varnothing$ (see Section 3.1).

We can combine several path sampling techniques into an MMIS estimator. For T marginal techniques p_i, and n_i technique-sample pairs $(\bar{\mathbf{x}}, \bar{\mathbf{t}}) \sim p_i(\bar{\mathbf{x}}, \bar{\mathbf{t}})$, where $\bar{\mathbf{t}} \in \mathcal{T}_i$, the estimator reads

$$\langle I_k \rangle_{\mathrm{MPS}} = \sum_{i=1}^{T}\sum_{j=1}^{n_i} \frac{f(\bar{\mathbf{x}}_{i,j})}{\sum_{i'=1}^{T}\sum_{j'=1}^{n_{i'}} p_{i'}(\bar{\mathbf{x}}_{i,j}|\bar{\mathbf{t}}_{i',j'})} \,, \tag{13}$$

which we refer to as our *marginal path sampling* (MPS) estimator.

Note that the MPS estimator (13) assumes the conditional density $p_{i'}(\bar{\mathbf{x}}_{i,j}|\bar{\mathbf{t}}_{i',j'})$ is readily computable. This is usually true for path sampling techniques, as we demonstrate in the following sections. Further, note that the auxiliary variables $\bar{\mathbf{t}}_{i,j}$ can only come from, and are used with, the associated marginal technique. These variables do not necessarily need to (all) be vertices; for example, they can include a randomly sampled wavelength, when doing spectral rendering.

5 MULTI-VERTEX PATH FILTERING

Path filtering [Keller et al. 2014; West et al. 2020] refines the contributions of sampled paths by connecting the prefix of each path to the suffixes of other paths (e.g., associated with different pixels).

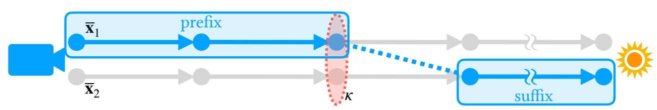

West et al. [2020] proposed an MIS interpretation of this method, over a continuous space of sampling techniques, where all possible prefixes within the filtering kernel around a vertex are techniques, and their corresponding suffixes are the samples. Their balance-heuristic continuous MIS (CMIS) formulation involves a marginal density of suffix samples which is an intractable integral over prefixes. To that end, West et al. applied stochastic MIS (SMIS) to approximate that marginal using the finite number of sampled prefixes.

Under our marginal path-sampling framework (13), we can view a prefix technique as a set of auxiliary vertices $\bar{\mathbf{t}}$, and the constructed path as a set of sample vertices $\bar{\mathbf{x}}$.

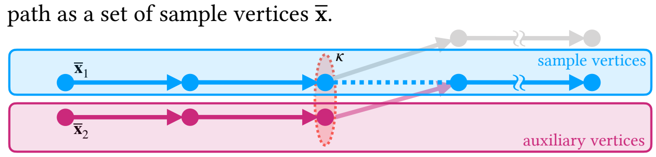

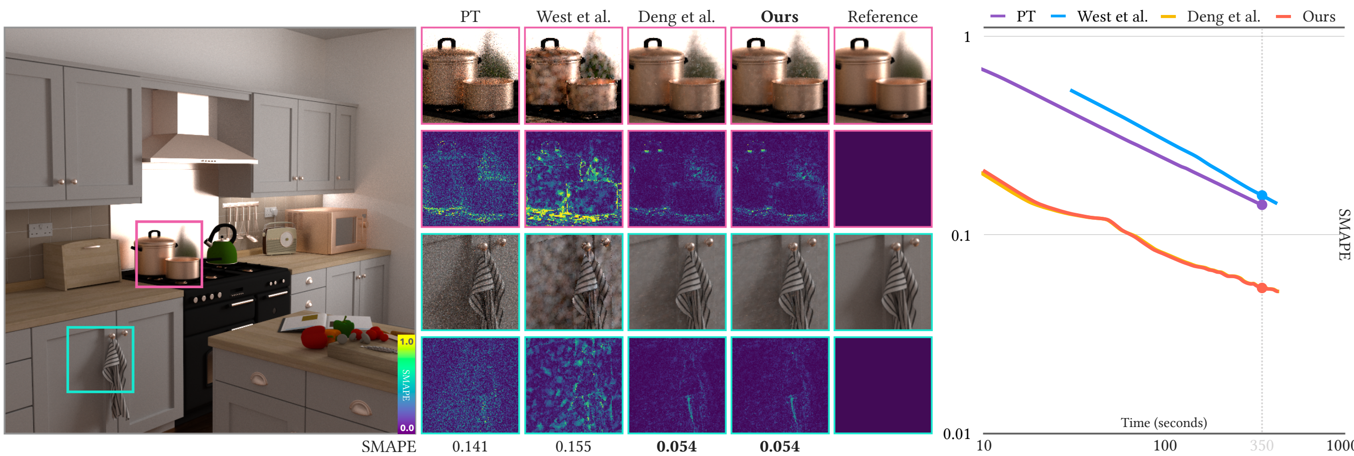

Fig. 2. Comparison of several methods after 350s of rendering. Path tracing (PT) renders every pixel independently and shows noise. West et al.'s [2020] filtering reuses vertices to construct novel paths but its computationally intensive weight computation affords tracing only a small number of initial paths whose contributions are blurred into conspicuous artifacts. Deng et al.'s [2021] iterative filtering and our multi-vertex filtering maximize vertex reuse and are much faster, showing significant improvement. Though different in formulation, these two methods are similar in implementation and performance.

One or more prefixes may satisfy the filter kernel κ at a query point. We have one classical technique—the suffix sampled from the query-point vertex—and zero or more marginal techniques (i.e., connections to other suffixes). These can be combined into an MPS estimator (13).

Extension to multi-vertex filtering. Filtering is traditionally performed at a single select vertex of each path [Keller et al. 2014; West et al. 2020]. Our framework allows us to maximize the amortization of path sampling by filtering at multiple vertices along the path.

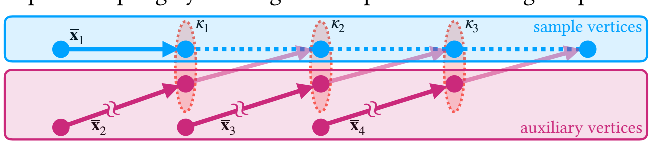

In the inline figure above, the last three (blue) output-path vertices are each collected from a different prefix. The set of (purple) auxiliary vertices represents the corresponding conditioning prefixes. Our set of conditional techniques is then all of the ways we can sample the constructed path—the choices of prefix vertices that fall within the support of filtering kernel along the path. If we apply filtering at k vertices, there are $T = \prod_{i=1}^{k} N(\mathbf{P}, \kappa_i)$ conditional path sampling techniques , where $\mathbf{P}$ is the set of traced paths and $N(\mathbf{P}, \kappa_i)$ counts the number of prefixes that fall within the support. This results in a number of techniques exponential in k.

Computing the MMIS weights in the resulting MPS estimator requires evaluating the PDF of $\bar{\mathbf{x}}$ for each of the exponential number of conditional techniques. To this end, note that density of sampling each vertex reappears in multiple (conditional) path PDFs in the denominator of the estimator (13):

$$\sum_{i'=1}^{T} \sum_{j'=1}^{n_{i'}} p_{i'}(\bar{\mathbf{x}}_{i,j} | \bar{\mathbf{t}}_{i',j'}) = \sum_{i'=1}^{T} \sum_{j'=1}^{n_{i'}} \prod_{k'=1}^{k} p_{i'}(\mathbf{x}_{i,j,k'} | \bar{\mathbf{y}}_{i',j',k'}), \quad (14)$$

where, $\mathbf{x}_{i,j,k'}$ is the k'-th vertex from the sensor on the j-th path sample from the i-th marginal technique, and $\bar{\mathbf{y}}_{i',j',k'}$ is the prefix that conditions the PDF evaluation. The index i iterates over the permutations of prefix choices at each kernel, and thereby intrinsically parameterizes a choice of prefix at each kernel.

Equation (14), when expanded, has many repeated terms and can be reorganized from a sum of products into a product of sums. Such a reorganization reduces the computational complexity from exponential to linear, and is analogous to locally performing MIS over each filtering kernel and taking their product. We provide technical details in the supplemental document.

5.1 Discussion

Our path-filtering MPS estimator is unbiased when there is at least one sampling technique where $p_i(\bar{\mathbf{x}}|\bar{\mathbf{t}}) > 0$ when $f(\bar{\mathbf{x}}) > 0$ (see the supplemental document). Including the classical path tracing techniques (i.e., the natural suffix at the query vertex) guarantees this condition even if all marginal techniques are biased.

However, an unbiased implementation can be prohibitively expensive as it requires evaluating visibility and BSDFs along connections [West et al. 2020]; this problem is compounded when filtering at multiple vertices. In practice, we use a biased approximation that assumes all vertices within a kernel have the same visibility. This allows us to avoid additional ray casting, and to compute the weights and throughput terms locally at each filtering kernel and iteratively propagate them (see supplemental document for details). This iterative propagation of localized terms is similar to the method of Deng et al. [2021]. Additionally, both Deng et al.'s and our implementation perform the highly localized filtering operations on the GPU. Not only is path filtering well suited to the highly parallel processing of GPUs, we have found that a GPU implementation is practically essential for achieving satisfactory performance.

Though our biased implementation is analogous to that of Deng et al., there are three key differences. The first is in the formulation. Ours is an MMIS-based path-integral estimator, providing a clear connection between the filtering process and the underlying transport integrals. The fixed-point iteration of Deng et al., while

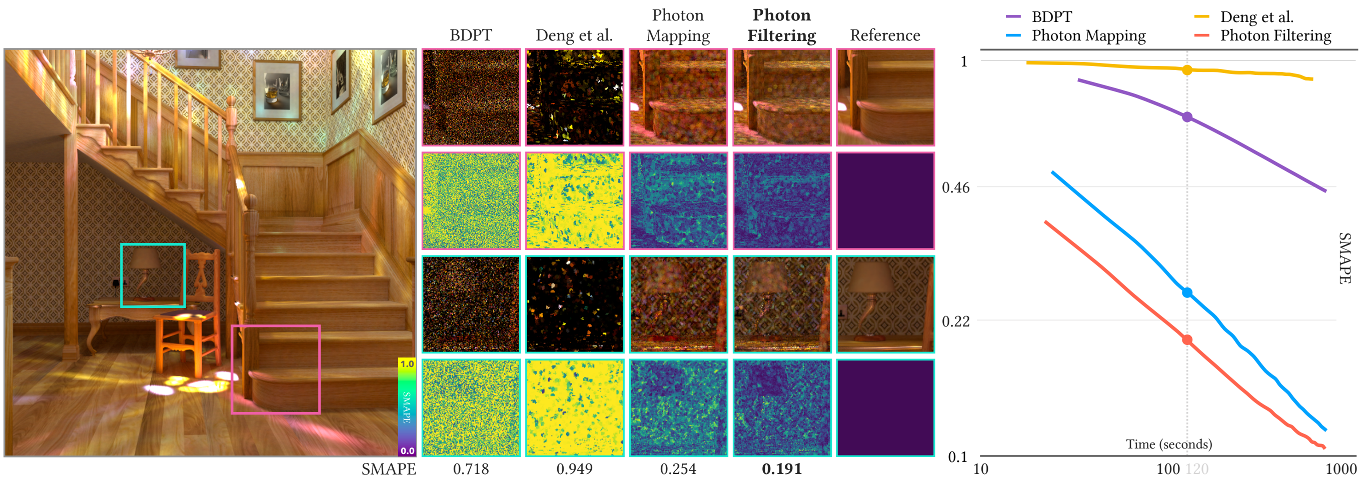

Fig. 3. A scene lit predominantly by caustics and indirect light from caustics, which is challenging for both bidirectional path tracing and the filtering method of Deng et al. [2021]. Photon mapping shows improvement, but exhibits artifacts in indirectly lit areas. Our filtered variant better handles the complex lighting in this scene and shows reduced error in indirectly lit areas. The zoom-ins show results after 120s of rendering, and all methods use a fixed-radius kernel.

numerically sound, lacks apparent connection to MC estimation. Our method does not need energy clamping, and our formulation enables other forms of multi-vertex filtering (e.g., for photon-density estimation, see Section 6).

The second difference is in the sample weighting. Our MMIS weights are principally based on the balance heuristic via marginals, which permits unbiased implementation in principle. The weights used by Deng et al. [2021] are not guaranteed to sum up to one.

The third difference is in clustering method. Existing path-filtering methods typically employ spatial nearest-neighbor lookup at each query vertex independently, which results in a quadratic number of PDF evaluations. Deng et al. [2021] use k-means clustering which greatly amortizes computational overhead but lacks controllable bias. We employ a clustering variant (described in the supplemental document) that introduces controllable bias by only clustering vertices that are within a radius parameter.

5.2 Results

We implemented the sampling part of our method in a CPU-based ray tracer and the filtering part in CUDA. Experiments were performed on consumer-grade hardware with a Ryzen 5 3600 CPU and an RTX 3070Ti (8GB) GPU.

Figure 2 shows a same-time comparison of various methods. Path tracing (PT) produces a noisy output as it renders every pixel independently. Single-vertex filtering [West et al. 2020] employs a computationally expensive weighting scheme which limits the number of initial paths that can be traced in the allocated time, whose contributions are blurred into conspicuous splotches. Deng et al.'s [2021] iterative filtering and our multi-vertex filtering appear significantly more converged, especially in mostly indirectly lit areas. While multi-vertex path filtering also correlates samples, the greatly lower error makes correlation artifacts much less noticeable.

6 MULTI-VERTEX PHOTON FILTERING

Path filtering demonstrates an appreciable reduction in estimation variance over unfiltered path tracing. However, both perform poorly under illumination that is poorly importance sampled from the camera, e.g., strongly directional lighting or caustics (see Fig. 3). Such scenes are better suited to photon-density estimation [Jensen 1996] where paths are sampled from emitters instead. Inspired by the approach from Section 5, we also propose an adjoint method that augments photon tracing with multi-vertex filtering.

Photon-density estimation can be interpreted as construction of a path by *merging* a sensor subpath $\overline{\mathbf{x}}_1$ with an emitter subpath $\overline{\mathbf{x}}_2$ [Georgiev et al. 2012], if the last vertex of $\overline{\mathbf{x}}_2$ lands within the support of a given kernel (κ_{de} in the illustration below) around the last vertex of $\overline{\mathbf{x}}_1$. The probability of a successful merge is proportional to the kernel extent (represented below by a cone).

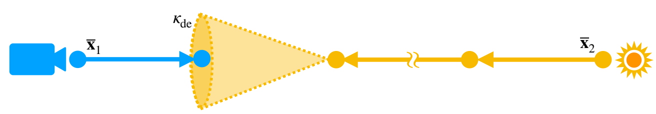

While photon-density estimation considers the merging of sensor and emitter subpaths, it is not uncommon for vertices of different emitter subpaths to land near each other. This provides an opportunity, much like in path filtering, to improve sample contribution by filtering the contributions of emitter subpaths. The distinction between sensor and emitter subpaths does not necessitate an entirely new method, since our MPS framework is based on general path sampling conditionally on a set of auxiliary vertices. We can use the same procedure as in Section 5, where the marginal sampling PDFs $p_i(\overline{\mathbf{x}})$ of Eq. (13) additionally include the probability of merging.

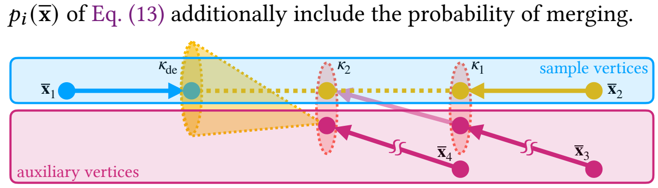

This adjoint application is not trivial within the frameworks of other filtering methods [Keller et al. 2014; Deng et al. 2021] as they are not based on path sampling. The possibility of performing filtering multiple times along a bidirectionally constructed path was also conceived by Hachisuka et al. [2012], but only in passing.

Multi-vertex photon filtering can be implemented as a post-process after emitter tracing, which operates as in path filtering but with the sampling directions reversed. Once the subpaths have been filtered, standard photon-density estimation is performed as usual.

6.1 Results

Our implementation is based on the same rendering system and hardware as in Section 5.2. The scene in Fig. 3 is predominantly lit by caustic illumination from the top which is exceptionally challenging for (bidirectional) path tracing methods. Conversely, it is exactly the type of scene where photon mapping excels. Deng et al.'s path sampling from the camera and exhibits high variance. Unfiltered photon mapping yields significantly lower error. Enabled by our marginal path sampling framework, our multi-vertex photon filtering reduces error further.

It is worth noting that our multi-vertex photon filtering does not improve the distribution or number of photons. Instead, it allows each emitter subpath to carry the contribution of multiple subpaths. This can be advantageous in closed scenes and those with strong indirect light contribution.

7 LIMITATIONS AND FUTURE WORK

Limitations. Our path sampling framework assumes that we can compute the PDF of a vertex given the auxiliary variables that condition its sampling. For any technique where the true distribution of samples is unknown, like those used in Markov chain Monte Carlo and adaptive sampling methods, we are unable to compute this conditional PDF.

Both path and photon filtering require additional memory to store information about the vertices of each path and their weights. This memory overhead can grow prohibitively large for high-resolution images. Deng et al. [2021] also mentioned that their method consumes 800 MB for 1280×720 images. This added storage cost may not be justifiable depending on the setting.

Future work. Path and photon filtering are each effective in certain lighting situations but can perform poorly in others. Rather than choosing one, ideally the two would be combined in a unified method like VCM/UPS [Georgiev et al. 2012; Hachisuka et al. 2012]. We believe our framework should in principle enable the derivation of such a unified method. Another interesting direction would be to devise a vertex-clustering scheme that would make the filtering estimator consistent as the cluster size is reduced (as in progressive photon mapping [Hachisuka et al. 2008]), or to perhaps fully debias the estimator [Misso et al. 2022].

Similarly to stochastic MIS [West et al. 2020], our marginal MIS is a stochastic approximation of an estimator that uses exact marginals. This approximation introduces additional variance which we show vanishes empirically as more techniques are sampled. Increasing the number of techniques, however, comes at the ost of additional computational overhead. A thorough variance analysis of marginal MIS may provide insight on the optimal number of techniques that balances variance reduction and computational overhead.

The applications shown in this paper are small examples of what is possible with our framework. One could consider reusing the same subpath *multiple times* to construct a complete path, essentially forming a graph cycle. Deng et al. [2021] state that cycles can be formed by iterating path filtering. Our framework also encompasses the idea of cycles, as path samples containing repeated vertex sequences. However, such repetitions introduce high correlation, which makes their practical variance-reduction potential unclear. Our framework could help identify how path-constructed graphs and cycles should be properly used to solve light-transport simulation.

Through preliminary experiments, we have also found that one might need to carefully select a subset of efficient sampling techniques out of a large number of available techniques. Our framework enables such sub-sampling and its exploration remains an open problem.

ACKNOWLEDGMENTS

This work has been partially funded by a grant from Autodesk and a Japanese Government (MEXT) Scholarship.

We express our gratitude to the anonymous reviewers for their valuable feedback, and to the following Blend Swap users for the rendering assets used in this paper: *Jay-Artist* for the kitchen scene in Fig. 2, and *Wig42* for the stairwell scene in Fig. 3.

REFERENCES

Pablo Bauszat, Victor Petitjean, and Elmar Eisemann. 2017. Gradient-Domain Path Reusing. *ACM Trans. Graph.* 36, 6, Article Article 229 (Nov. 2017), 9 pages. https://doi.org/10.1145/3130800.3130886

Philippe Bekaert, Mateu Sbert, and John Halton. 2002. Accelerating Path Tracing by Re-Using Paths. In *Proceedings of the 13th Eurographics Workshop on Rendering*. Eurographics Association, 125–134.

Xi Deng, Miloš Hašan, Nathan Carr, Zexiang Xu, and Steve Marschner. 2021. Path Graphs: Iterative Path Space Filtering. *ACM Trans. Graph.* 40, 6, Article 276 (dec 2021), 15 pages. https://doi.org/10.1145/3478513.3480547

Iliyan Georgiev, Jaroslav Křivánek, Tomáš Davidovič, and Philipp Slusallek. 2012. Light Transport Simulation with Vertex Connection and Merging. 31, 6 (Nov. 2012), 192:1–192:10. https://doi.org/10/gbb6q7

Toshiya Hachisuka, Shinji Ogaki, and Henrik Wann Jensen. 2008. Progressive Photon Mapping. 27, 5 (Dec. 2008), 130:1–130:8. https://doi.org/10/cn8h39

Toshiya Hachisuka, Jacopo Pantaleoni, and Henrik Wann Jensen. 2012. A Path Space Extension for Robust Light Transport Simulation. 31, 6 (Jan. 2012), 191:1–191:10. https://doi.org/10/gbb6n3

Henrik Wann Jensen. 1996. *The Photon Map in Global Illumination.* Ph.D. Thesis. Technical University of Denmark.

James T. Kajiya. 1986. The Rendering Equation. 20, 4 (Aug. 1986), 143–150. https://doi.org/10/cvf53j

Alexander Keller, Ken Dahm, and Nikolaus Binder. 2014. Path Space Filtering *(SIGGRAPH '14)*. ACM, 68:1–68:1. https://doi.org/10/gfz6mr

Zackary Misso, Benedikt Bitterli, Iliyan Georgiev, and Wojciech Jarosz. 2022. Unbiased and consistent rendering using biased estimators. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 41, 4 (2022). https://doi.org/10.1145/3528223.3530160

Hao Qin, Xin Sun, Qiming Hou, Baining Guo, and Kun Zhou. 2015. Unbiased Photon Gathering for Light Transport Simulation. *ACM Trans. Graph.* 34, 6, Article 208 (Oct. 2015), 14 pages. https://doi.org/10.1145/2816795.2818119

Eric Veach. 1997. *Robust Monte Carlo Methods for Light Transport Simulation.* Ph.D. Thesis. Stanford University, United States – California.

Eric Veach and Leonidas J. Guibas. 1995. Optimally Combining Sampling Techniques for Monte Carlo Rendering, Vol. 29. 419–428. https://doi.org/10/d7b6n4

Rex West, Iliyan Georgiev, Adrien Gruson, and Toshiya Hachisuka. 2020. Continuous Multiple Importance Sampling. *ACM Transactions on Graphics (Proceedings of SIGGRAPH)* 39, 4 (July 2020). https://doi.org/10.1145/3386569.3392436