Robust Light Transport Simulation via Metropolised Bidirectional Estimators

Martin Šik[1]* Hisanari Otsu[2] Toshiya Hachisuka[2] Jaroslav Křivánek[1]
[1]Charles University in Prague [2]The University of Tokyo

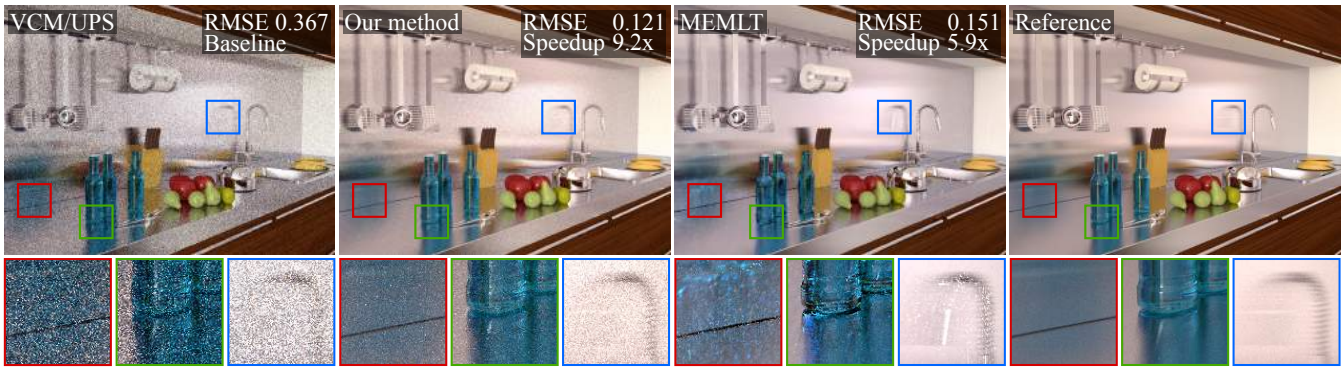

Figure 1: *We present a robust light transport simulation algorithm that fuses key ideas behind vertex connection and merging/unified path sampling (VCM/UPS) [Georgiev et al. 2012; Hachisuka et al. 2012] and Markov chain Monte Carlo methods. It outperforms state-of-the-art approaches, such as VCM/UPS or Metropolis light transport with manifold exploration (MEMLT) [Jakob and Marschner 2012], especially in scenes with complex glossy or specular transport and complex visibility. At the same time, the algorithm is substantially simpler than many of the recent Metropolis light transport variants.*

Abstract

Efficiently simulating light transport in various scenes with a single algorithm is a difficult and important problem in computer graphics. Two major issues have been shown to hinder the efficiency of the existing solutions: light transport due to multiple highly glossy or specular interactions, and scenes with complex visibility between the camera and light sources. While recent bidirectional path sampling methods such as vertex connection and merging/unified path sampling (VCM/UPS) efficiently deal with highly glossy or specular transport, they tend to perform poorly in scenes with complex visibility. On the other hand, Markov chain Monte Carlo (MCMC) methods have been able to show some excellent results in scenes with complex visibility, but they behave unpredictably in scenes with glossy or specular surfaces due to their fundamental issue of sample correlation. In this paper, we show how to fuse the underlying key ideas behind VCM/UPS and MCMC into a single, efficient light transport solution. Our algorithm is specifically designed to retain the advantages of both approaches, while alleviating their limitations. Our experiments show that the algorithm can efficiently render scenes with both highly glossy or specular materials *and* complex visibility, without compromising the performance in simpler cases.

Keywords: Global illumination, light transport simulation, Markov chain Monte Carlo, photon density estimation

Concepts: •**Computing methodologies** → **Ray tracing;**

*sik@cgg.mff.cuni.cz

SA '16 Technical Papers,, December 05 - 08, 2016, , Macao

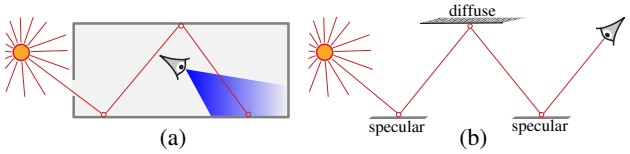

Figure 2: *Two major issues that compromise efficiency of current light transport algorithms: (a)* Complex visibility *– the visible area (blue) is lit only through multiple reflections and/or much of the light emitted from the sources never reach the visible area. (b)* Complex glossy and specular transport *– light paths with chains of interactions on highly glossy (or specular) surfaces (e.g. a specular-diffuse-specular path shown in this figure).*

1 Introduction

Monte Carlo light transport simulation has become a de-facto standard tool for photorealistic rendering. While the existing simulation algorithms often work well, two major issues have been shown to still compromise the efficiency of many algorithms: light transport due to multiple highly glossy or specular interactions, and scenes with complex visibility between the camera and light sources (Fig. 2). Carefully avoiding such cases by modifying the scene can be cumbersome because to do that, the users need to understand how the algorithms work and when they may fail. Developing robust solutions that work well under various scene configurations has thus received a well-deserved attention.

One widely accepted idea for dealing with glossy and specular transport is the use of *bidirectional path estimators*. The key observation here is that, given a full light transport path between a light source and the camera, one can construct this path in various ways, each of which is efficient at capturing different lighting effects. Examples of bidirectional estimators include bidirectional path tracing [Veach and Guibas 1994], photon density esti-

ISBN: 978-1-4503-4514-9/16/12
DOI: http://dx.doi.org/10.1145/2980179.2982411

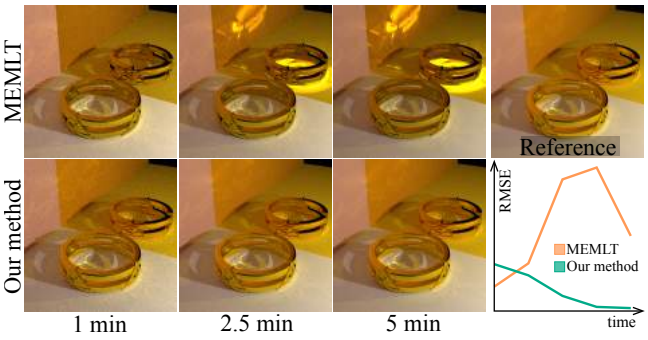

Figure 3: *Metropolis light transport with Manifold exploration (MEMLT) [Jakob and Marschner 2012] suffers from excessive sample correlation in scenes with specular surfaces, which results in its irregular, unpredictable convergence (top row). Despite being based on MCMC sampling, our method (bottom row) shows good convergence properties typical of regular Monte Carlo algorithms: unconverged results provide a reliable preview of the final image, and are amenable to denoising.*

mation [Jensen 1996], and combinations of the two [Georgiev et al. 2012; Hachisuka et al. 2012; Křivánek et al. 2014a]. The latter works have shown that a key to the good efficiency of photon density estimation is an extensive *subpath reuse*. Bidirectional estimators themselves, however, do not perform well in scenes with complex visibility due to low probability of combining two subpaths [Popov et al. 2015].

Concerning complex visibility, Markov chain Monte Carlo (MCMC) algorithms have been able to deliver some excellent results. The key idea is to sequentially sample a new light transport path based on the previous one, which then forms a *Markov chain* of paths. Since the chain often visits important and unoccluded paths, this class of algorithms naturally deals with complex visibility [Veach and Guibas 1997; Hachisuka and Jensen 2011; Jakob and Marschner 2012; Kaplanyan and Dachsbacher 2013; Hachisuka et al. 2014]. The MCMC algorithms are, however, also known to behave unpredictably in scenes with multiple glossy and specular interactions. Some recent works have indeed highlighted their unpredictable convergence behavior and various image artifacts due to the fundamental issue of sample correlation in MCMC [Křivánek et al. 2014b; Vorba et al. 2014]. Fig. 3 illustrates the problem and shows that our algorithm is, to a large extent, immune to this issue.

We propose a practical light transport simulation algorithm that fuses the two key concepts mentioned above into a single solution. Our algorithm is designed to gracefully handle scenes both with complex glossy and specular transport as well as complex visibility. To achieve this goal, we take advantage of bidirectional estimators and subpath reuse to deal with the glossy/specular transport, while utilizing Markov chain Monte Carlo to address complex visibility. While it is natural to expect that fusing these ideas could work well, a combination of bidirectional estimators and MCMC in fact still leaves a large design space of different approaches. The first proposal of a practical and efficient combination, based on a number of algorithmic considerations, is our key contribution.

To be concrete, the final algorithm is built upon the vertex connection and merging (VCM) [Georgiev et al. 2012], a.k.a. unified path space sampling (UPS) [Hachisuka et al. 2012] algorithms (referred to as VCM/UPS from now on). To avoid inheriting the issues of correlation and unpredictable convergence of MCMC, we propose to use a mixture of two target functions: one based solely on path visibility and the other on the path contribution. Doing so makes the target distribution easier to explore by a Markov chain. Furthermore, we reduce correlation in the image space and improve image-

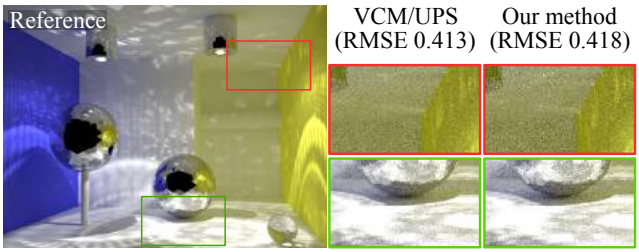

Figure 4: *Our algorithm does not significantly impair the performance of VCM/UPS in simple scenes, where most of the light subpaths generated by independent Monte Carlo contribute to the image. We ran both methods for 5 minutes to render these images.*

plane stratification by decoupling the sampling of light subpaths (i.e. subpaths from a light source) and eye subpaths (i.e. subpaths from the camera), and by sampling the eye subpaths by ordinary (independent) Monte Carlo.

The results demonstrate that our algorithm is indeed robust with respect to various scene configurations, see Fig. 1 for an example. An important practical aspect is that our algorithm has a low overhead over baseline VCM/UPS; as such it exhibits good performance even in relatively simple cases, as we show in Fig. 4.

2 Related Work

Bidirectional methods. Robustness in rendering refers to acceptable efficiency for any practical input scene [Veach 1997]. Veach and Guibas [1994] developed bidirectional path tracing (BPT), concurrently with Lafortune and Willems [1993], as a more robust alternative to simple unidirectional algorithms. The Veach and Guibas's variant [1995] is more efficient thanks to the use of multiple importance sampling (MIS) for combining the various bidirectional estimators; our algorithm also relies on MIS.

A different line of research was started by Jensen [1996]; his photon mapping has became widely used due to its simplicity and ability to efficiently render the so-called specular-diffuse-specular (SDS) paths (e.g. indirectly visible caustics) as illustrated in Fig. 2. The progressive formulation of photon mapping due to Hachisuka et al. [2008] has subsequently resulted in a renewed interest in photon density estimation [Hachisuka and Jensen 2009; Knaus and Zwicker 2011; Jarosz et al. 2011]. Recently, Georgiev et al. [2012] and Hachisuka et al. [2012] have independently proposed a reformulation of photon density estimation and BPT that enables their combination into an algorithm that we refer to as VCM/UPS. Křivánek et al. [2014a] later extended these ideas for participating media. While VCM/UPS is robust with respect to various lighting conditions including diffuse indirect illumination and SDS paths, its performance can degrade with complex visibility. In such cases, the bidirectional estimators in VCM/UPS often fail to sample paths with a non-zero contribution because light subpath tracing is unaware of the camera location [Popov et al. 2015]. Our algorithm deals with this issue by using Markov chain Monte Carlo to generate light subpaths that make non-zero contributions to the image.

Markov chain Monte Carlo (MCMC). Veach and Guibas [1997] proposed Metropolis light transport (MLT) as an application of Metropolis-Hastings (MH) sampling to path integral estimation. The MH algorithm belongs to a wider family of Markov chain Monte Carlo methods [Brooks et al. 2011]. Pauly et al. [2000] extended MLT to handle participating media. Recent works have focused on improving local exploration in MLT via mutation strategies that take the differential geometry of the path space into account [Jakob and Marschner 2012; Hanika et al. 2015; Li et al. 2015]. Kelemen et al. [2002] proposed a primary sample space

formulation of MLT, where MCMC sampling manipulates a set of "random" numbers that define light transport paths. Multiplexed MLT [Hachisuka et al. 2014] reformulates the primary sample space MLT such that suitable path sampling techniques are more frequently selected by the MCMC sampler.

A fundamental advantage of MCMC algorithms is that a path which would be blocked (i.e. invisible) is never accepted as a valid sample. This makes MCMC an excellent approach to deal with complex visibility. At the same time, though theoretically sound, the use of MCMC in light transport simulation results in some major issues. Firstly, MCMC algorithms lose explicit control of stratification over the image plane. Energy redistribution path tracing [Cline et al. 2005] strives to retain stratification by running many short Markov chains, but the underlying algorithm is still equivalent to MLT (i.e. the used mutations and the target functions are the same) and so the success has been only partial. Secondly, the path contribution function, which often serves as the target distribution of the MCMC sampler, has many peaks and discontinuities. This situation leads to strong correlation artifacts due to high rejection rates and to unpredictable convergence behavior due to a chain being stuck at one sample for a long time. The current MCMC algorithms are thus not suitable for fast pre-visualization of the final images [Hoberock and Hart 2010; Křivánek et al. 2014b]. Furthermore, unconverged images generated by MCMC algorithms are not amenable to denoising [Zwicker et al. 2015]. An essential aspect of our algorithm design is to substantially reduce those problems of MCMC rendering algorithms.

Several works propose to use MCMC methods for sampling *sub*paths, as opposed to full light transport paths. Fan et al. [2005] run the MLT algorithm and store selected path vertices to be used for photon density estimation. Segovia et al. [2007] use a similar idea for sampling virtual point lights [Dachsbacher et al. 2014]. Hachisuka and Jensen [2011] propose a photon tracing algorithm where the target function of the MCMC sampler is defined as a binary path visibility function. This function equals to one for light subpaths that make a contribution to any of a pre-generated set of eye subpaths by photon density estimators. Such a simple target function is much easier to explore by the MCMC sampler, which, in turn, reduces the correlation problems mentioned above. We follow a similar idea, but our path visibility takes into account all the VCM/UPS path estimators, not just photon density estimators. The algorithm of Chen et al. [2011] bears similarity to that of Hachisuka and Jensen, but uses a somewhat arbitrary target function instead of path visibility.

Replica exchange [Swendsen and Wang 1986] is another important idea to reduce correlation in MCMC sampling that has been previously applied in light transport [Kitaoka et al. 2009; Hachisuka and Jensen 2011]. We also employ replica exchange to further improve the path space exploration in our MCMC sampler.

Path guiding and probabilistic connections. Apart from MCMC methods, another option to deal with complex visibility is importance-driven sampling of light emission and scattering directions, also referred to as *path guiding* [Jensen 1995; Bashford-Rogers et al. 2012]. Vorba et al. [2014] have shown that the use of path guiding in VCM/UPS can significantly improve its performance. Another way to handle complex visibility is to importance sample subpaths for connections [Popov et al. 2015] in order to increase a chance of non-zero contribution. Since path guiding and importance sampling of connections is orthogonal to MCMC under the primary sample space formulation, both methods could potentially be combined with our method. As our focus is on exploring the improvement thanks to MCMC over baseline VCM/UPS, we do not utilize path guiding or importance sampling of connections in our experiments.

3 Background

In order to design an algorithm that combines the strengths of bidirectional estimators with path reuse and MCMC sampling, we first give an overview of the basic elements that these methods build on.

3.1 Bidirectional sampling and VCM/UPS

Path integral formulation. Light transport simulation can be expressed by the path integral $I_j = \int_\Omega h_j(\bar{x}) f(\bar{x}) d\mu(\bar{x})$, where I_j is the pixel value of the j-th pixel, Ω is the space of all possible paths, $\bar{x} \in \Omega$ is a full transport path from a light source to the camera. A path of length k can be represented as a vector of vertices $\bar{x} = (x_0, \ldots, x_k)$. The path contribution function f gives the amount of light energy transported along the path $\bar{x}$. h_j is the pixel filter of the pixel j [Veach 1997].

Given a *path sampling technique* (e.g. path tracing or light tracing) that generates N random paths $\bar{x}_i$ according to the probability density function (pdf) $p(\bar{x}_i)$, we can estimate the path integral using a Monte Carlo estimator $\langle I_j \rangle = \frac{1}{N} \sum_{i=1}^{N} \frac{h_j(\bar{x}_i) f(\bar{x}_i)}{p(\bar{x}_i)}$.

Bidirectional path tracing. Bidirectional path tracing (BPT) employs a family of different path sampling techniques and combines their results using multiple importance sampling (MIS) [Veach and Guibas 1995]. The algorithm repeatedly generates a pair of subpaths: an eye subpath from the camera and a light subpath from a light source. The two subpaths are then connected to form a full path. The various bidirectional sampling techniques used in BPT differ by the lengths of the respective eye and light subpaths, and each is effective at sampling different lighting effects.

VCM/UPS. The VCM/UPS algorithm [Georgiev et al. 2012; Hachisuka et al. 2012] improves upon BPT by adding a new family of sampling techniques based on photon density estimation [Jensen 1996]. Georgiev et al. termed them *merging*, as they can be thought as welding two spatially close vertices (from light and eye subpaths) into one, thereby introducing *regularization* [Kaplanyan and Dachsbacher 2013]. The VCM/UPS algorithm shows that the efficiency of merging techniques relies on *path reuse*: *all* the traced subpaths are reused by *each* subpath in the opposite direction.

3.2 Markov Chain Monte Carlo in Light Transport

Markov chain Monte Carlo (MCMC) methods can generate samples according to any target function $p^*(\bar{x}) \propto p(\bar{x})$ without knowing the normalization constant $b = \int_\Omega p^*(\bar{x}) d\bar{x}$. This can be used to generate paths proportionally to the path contribution function $f(\bar{x})$, which, in turn, should minimize the variance of the resulting path integral estimators.

The Metropolis-Hastings algorithm [Hastings 1970] is a common MCMC method that generates samples as a history of a Markov chain. Given a current state $\bar{x}_i$, a candidate state $\bar{y}$ is proposed using a *mutation* according to a *proposal distribution* $q(\bar{x}_i \to \bar{y})$. The new state is then set as $\bar{x}_{i+1} = \bar{y}$ with the *acceptance probability*

$$a = \min\left\{1, \frac{p^*(\bar{y}) q(\bar{y} \to \bar{x}_i)}{p^*(\bar{x}_i) q(\bar{x}_i \to \bar{y})}\right\}. \quad (1)$$

Otherwise, the chain repeats the current state, $\bar{x}_{i+1} = \bar{x}_i$. In Metropolis Light Transport (MLT) [Veach and Guibas 1997], each state of the Markov chain corresponds to a full transport path, and mutations operate directly in the path space.

Primary sample space MLT (PSSMLT). Our algorithm builds upon primary sample space MLT (PSSMLT) [Kelemen et al. 2002]. PSSMLT simplifies the original MLT by applying mutations on

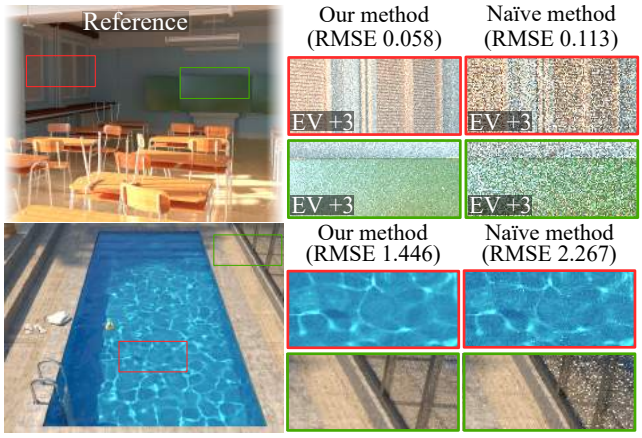

Figure 5: *Equal-time comparison (1 hour) between our method and a naïve combination of PSSMLT and VCM/UPS with photons generated by visibility-driven MCMC [Hachisuka and Jensen 2011]. EV gives the exposure value adjustment used for the insets.*

a vector of random numbers (primary space samples) $u \in \mathcal{U}$, rather than directly on paths. The vector u is mapped to a set of paths through all the path sampling techniques from BPT, while using the vector u as the "random numbers" for path generation.

Replica exchange. If the target function has peaks and discontinuities, a MCMC sampler may fail to evenly explore the state space. This results in uneven convergence or even missing out entire state space regions. The replica exchange (a.k.a. parallel tempering) method alleviates such issues by using several interacting Markov chains [Swendsen and Wang 1986; Kitaoka et al. 2009; Hachisuka and Jensen 2011]. A common design is to have a main chain with the original target function and other chains with flatter target functions that allow easier exploration. The chains interact by swapping their current states; in this way one chain can help another explore a different part of the state space. The chains generate samples according to their target function if and only if the swaps between any two chains are performed with probability

$$r = \min\left\{1, \frac{p_1^*(\bar{x}_2)p_2^*(\bar{x}_1)}{p_1^*(\bar{x}_1)p_2^*(\bar{x}_2)}\right\}. \quad (2)$$

Here p_1^* and p_2^* are the target functions of the two chains, and $\bar{x}_1$ and $\bar{x}_2$ are their respective current states.

4 Algorithm Design

In this section, we summarize the design decisions that led to our algorithm and briefly mention possible alternatives.

4.1 Requirements

Our goal is to develop an algorithm that can efficiently render scenes with various kinds of (indirect) illumination, including *complex transport on glossy/specular surfaces*. At the same time, the algorithm should maintain good efficiency in scenes with *complex visibility* between light sources and the camera, such as natural light illuminating an interior through a small glass window pane. For the algorithm to be practical, it must *not impair the efficiency* of current state-of-the-art approaches *in the simple cases* where they already work well. And it should also be *relatively easy to implement* so as to facilitate its adoption in practice.

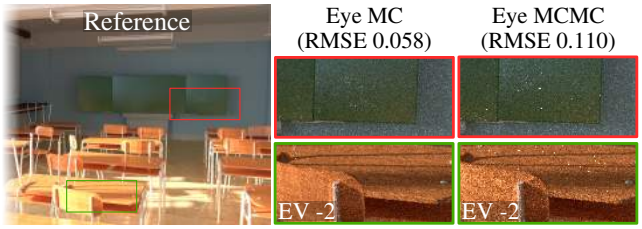

Figure 6: *Comparison of eye subpath generation by stratified Monte Carlo sampling (left) and by a visibility-driven MCMC sampler (right) in our algorithm. Eye subpath generation via MCMC is outperformed by stratified sampling in this, as well as in other scenes. EV gives the exposure value adjustment used for the insets.*

4.2 General Algorithm Design Strategy

We design the new algorithm based on ideas that have individually been shown to meet some of the above requirements. These include *bidirectional path estimators* and their *combination via multiple importance sampling (MIS)*, the vertex merging technique (i.e. photon density estimation) with its inherent *regularization* and large-scale *subpath reuse*. These ideas are the building blocks of VCM/UPS, which excels at handling complex glossy/specular transport, unlike simpler solutions such as path tracing. Furthermore, because VCM/UPS has been implemented in many popular rendering software solutions and our final algorithm is a lightweight extension of it, using VCM/UPS as a baseline meets our algorithmic simplicity criterion. The important cornerstone of our design is *MCMC path sampling*, which has complementary properties to VCM/UPS with respect to handling complex materials and visibility.

We pay attention not to inherit the shortcomings of VCM/UPS and MCMC sampling: bad performance of VCM/UPS in scenes with complex visibility due to light subpaths not being able to reach visually important regions; image artifacts and bad convergence behavior of MCMC due to excessive sample correlation; and insufficient image plane stratification in many existing MCMC approaches.

4.3 Design Choices

We choose VCM/UPS as the basis of our solution while using a MCMC sampler for path sampling. This high-level idea, however, leaves a number of unanswered questions: **Q1**. For which paths (eye, light, or full paths) should the MCMC sampler be used? **Q2**. What specific MCMC sampler should be used and what state space should it operate in (path space, primary sample space, or other)? **Q3**: How should its target function be defined? The specific design choices we make are directly motivated by the requirements stated above.

Fig. 5 illustrates the importance of finding suitable answers to these questions. We compare our method to a naïve algorithm that extends PSSMLT [Kelemen et al. 2002] by including all the sampling techniques from VCM/UPS. Light and eye subpaths are generated and combined as in original PSSMLT, but we use an additional set of light subpaths pre-generated by visibility-driven photon tracing of Hachisuka and Jensen [2011] to enable the merging technique. The poor performance of this algorithm compared to our solution shows that even within the relatively narrowed-down scope of combining VCM/UPS and MCMC sampling, the design choices can make a big difference in the resulting algorithm's performance (Sec.4 of the supplemental material discusses other possible combinations that we have considered in our research).

Q1: MCMC sampling of eye, light, or full paths. Previous work [Hachisuka and Jensen 2011; Segovia et al. 2007] has shown that using MCMC sampling to steer *light subpaths* towards visually

Figure 7: *Equal-sample comparison (5 samples per pixel) of rendering uniformly lit rectangle. Ideally, the resulting image should have uniform color. While MLT uses lens mutations to improve its image-plane stratification, it still delivers the worst results due to high sample correlation. While PSSMLT produces a slightly better result than MLT, our algorithm shows improvement over both MLT and PSSMLT thanks to our choice of using stratified eye subpaths.*

important regions can have tremendous efficiency benefits; we follow this idea. Likewise, one could expect that applying the same idea to *eye subpaths* in a combined algorithm – which is essentially symmetric with respect to the camera and light sources – should yield a similar benefit. However, we found the contrary to be true, as shown in Fig. 6. We explain this by the fact that high correlation of eye subpaths due to MCMC sampling compromises one of our design goals – good image plane stratification. For this reason, we opt for the use of regular, independent Monte Carlo to sample the eye subpaths. For the same reason, we also rule out the option of using MCMC sampling for *full* paths as in MLT. Fig. 7 illustrates image plane stratification/sample correlation of different methods.

Q2: MCMC sampling space. Motivated by the goal of algorithmic simplicity, we choose as our base MCMC algorithm a Metropolis-Hastings sampler (enhanced via replica exchange), that operates in the primary sample space. MCMC sampling directly in the path space would unnecessarily complicate the algorithm and go against one of our design principles. Furthermore, transformation from the primary sample space to path space naturally importance samples BRDFs and geometry terms along the path. Achieving this in MLT framework requires a rather complex transformation of paths from the path space to yet another space [Hanika et al. 2015].

Q3: Target function. In designing the target function, we face two contradicting goals: On the one hand, a target function with a simple shape (such as path visibility) is easily explored by a MCMC sampler, which reduces sample correlation and its negative practical consequences. On the other hand, a target function more similar to the actual path contribution can potentially reduce variance as per the importance sampling principle; but it may also lead to excessive sample correlation. Given the complementary pros and cons of the two options, we use two Markov chains, one with *path contribution* and the other with *path visibility* as their target functions. We refer to the two chains, respectively, as the *contribution chain* and the *visibility chain*. We let the two chains interact via replica exchange and we combine the samples produced by the two chains via MIS. This design contributes to our goal of reduced sample correlation in MCMC. Fig. 8 illustrates the complementary advantages of the two target functions and the benefit of their combination.

5 Algorithm

This section provides a complete description of our algorithm, starting with an overview shown in Algorithm 1. The algorithm works in iterations, each of which starts by generating all eye subpaths by an independent Monte Carlo sampler (line 2). We store all their vertices so that the merging technique can later be evaluated during the MCMC light subpath sampling. The eye subpaths are used to immediately evaluate the contribution of unidirectional path tracing and path tracing with next event estimation, as these sampling techniques are independent of light subpaths (line 3). We then execute our MCMC sampler, which is based on PSSMLT, to generate light subpaths and evaluate the remaining VCM/UPS techniques (line 4).

Algorithm 1 High-level overview of our algorithm.

1: **for** $i = 1 \dots M$ **do**
2: $\quad \mathcal{E} :=$ MCEYESUBPATHSAMPLING(N)
3: $\quad$ image += EVALUATEPATHTRACING($\mathcal{E}$)
4: $\quad$ image += MCMCLIGHTSUBPATHSAMPLING($\mathcal{E}$)
5: **end for**

5.1 Light Subpath Contribution

Contribution definitions. Let us start by formally defining the *light subpath contribution* that is used to evaluate the target function for both the contribution chain and the visibility chain. A single light subpath $\bar{x}_{\text{light}}$ in the VCM/UPS algorithm gives rise to an entire family of full paths created, respectively, by light tracing, bidirectional connections, and vertex merging [Georgiev et al. 2012]. We thus have multiple full paths generated by a single light subpath. We define the contribution of $\bar{x}_{\text{light}}$ as the sum over the individual MIS-weighted contributions of all these full paths. Let $\mathcal{P}(\bar{x}_{\text{light}})$ denote the set of such paths with the corresponding path sampling techniques that created them. The contribution of $\bar{x}_{\text{light}}$ is then

$$C(\bar{x}_{\text{light}}) = \sum_{(\bar{x},t)\in\mathcal{P}(\bar{x}_{\text{light}})} \frac{f(\bar{x})w_t(\bar{x})}{p_t(\bar{x})}. \tag{3}$$

Here $\bar{x}$ is a full path, t is a VCM/UPS sampling technique with the pdf p_t, f is the standard path contribution function (see Sec. 3.1), and $w_t(\bar{x})$ is our adjusted VCM/UPS multiple importance sampling (MIS) weight described below (Sec. 5.4). $C(\bar{x}_{\text{light}})$ is designed to measure the *total* image contribution due to the subpath $\bar{x}_{\text{light}}$. Note that since the contribution of the path tracing techniques is independent of $\bar{x}_{\text{light}}$, the paths created by these techniques are not in $\mathcal{P}(\bar{x}_{\text{light}})$.

Two target functions. We now use the light subpath contribution $C(\bar{x}_{\text{light}})$ to define the target function for the contribution and visibility Markov chains, respectively, as

$$p^*_{\text{con}}(C(\bar{x}_{\text{light}})) = \text{lum}(C(\bar{x}_{\text{light}})) \tag{4}$$

$$p^*_{\text{vis}}(C(\bar{x}_{\text{light}})) = \begin{cases} 1 & \text{if } C(\bar{x}_{\text{light}}) > 0 \\ 0 & \text{otherwise.} \end{cases} \tag{5}$$

That is, the former is given by the scalar luminance of the subpath contribution, while the latter is a binary indicator equal to one for light subpaths that make a non-zero contribution to the image.

Computing contribution. Contribution of a light subpath $C(\bar{x}_{\text{light}})$ consists of the contributions by bidirectional connections and by vertex merging. The merging contribution is computed using all pre-generated eye subpaths $\mathcal{E}$. In order to compute bidirectional connections we pick (at random) one eye subpath from $\mathcal{E}$ per each light subpath. As we already said, we exclude the contribution by path tracing techniques from $C(\bar{x}_{\text{light}})$. Fig. 9 shows how we combine the contributions from light and eye subpaths.

5.2 MCMC Sampler Details

We now provide details of our MCMC sampler used to generate light subpaths.

Initialization. To initialize our Markov chains, we apply the start-up bias elimination method [Veach and Guibas 1997]. The method draws an initial state using importance resampling from a small set (10000) of independent light subpaths.

Mutations. We use primary sample space MCMC, where we mutate a primary space sample $u \in \mathcal{U}$ using an adaptive mutation kernel (see the supplemental material for more details). Using u

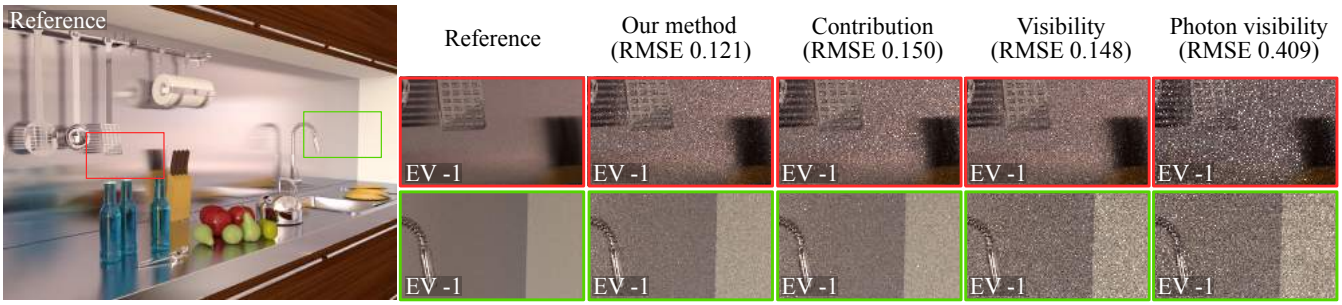

Figure 8: *Comparison of different target functions in our algorithm. From right to left, a target function based solely on photon visibility [Hachisuka and Jensen 2011] often yields suboptimal results. Our path visibility definition, which takes into account all the sampling techniques in VCM/UPS, considerably improves the results, but it may sometimes be outperformed by the target function given by the path contribution itself. Best overall results are achieved by our combined two-chain approach based on replica exchange.*

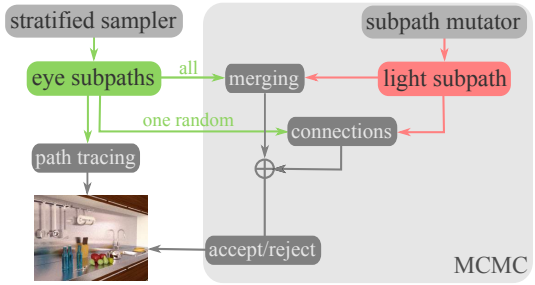

Figure 9: *Diagram showing how the contributions of light and eye subpaths are combined to compute the image. Each light subpath is first splatted onto all eye subpath vertices (merging) and then connected to a randomly selected eye subpath. The contributions due to both merging and connections are then combined via MIS weights and used for accepting/rejecting the light subpath as well as for computing the image. While the contribution by path tracing techniques is computed outside of the MCMC sampler, it also undergoes MIS weighting.*

as a vector of random numbers we then generate a light subpath $\bar{x}_{\text{light}}$ via local sampling techniques (i.e. light source emission and BRDF importance sampling). As in PSSMLT, we alternate between small steps (local mutations) and large steps (independent global mutations) to ensure good local and global exploration. However, the large steps are used only in the visibility chain, because replica exchange between the chains is sufficient to ensure global exploration in the contribution chain. In all scenes, we use the large step with probability $P(\text{largeStep}) = 0.3$.

Normalization. To estimate light transport using a MCMC sampler, we need the normalization constant $b_c = \int_{\mathcal{U}} p_c^*(C^*(u))du$ for both chains. Here, we have introduced the notation $C^*(u) = C(\bar{x}_{\text{light}})$, where $\bar{x}_{\text{light}}$ is the light subpath generated from u.

We compute the normalization using a regular Monte Carlo estimator, which utilizes the large steps of our MCMC algorithm. Given L large steps, the normalization is estimated as follows

$$\langle b_c \rangle = \frac{1}{L} \sum_{j=1}^{L} p_c^*(C(\bar{x}_{\text{light},j})). \tag{6}$$

We keep updating this estimate during the entire rendering process.

Combining the two chains. Using only samples from the contribution chain to render the image would be wasteful as half of the samples would be thrown away. We therefore combine the contributions of light subpaths generated by our two chains using another set of MIS weights (different from these used for VCM/UPS in Eq. (3)). The MCMC MIS weight for the given light subpath $\bar{x}_{\text{light}}$ from chain c is defined as

$$w_{\text{mcmc}}(c, \bar{x}_{\text{light}}) = \frac{\frac{p_c^*(C)}{b_c}}{\frac{p_{\text{vis}}^*(C)}{b_{\text{vis}}} + \frac{p_{\text{con}}^*(C)}{b_{\text{con}}}}, \tag{7}$$

where we have dropped the argument of $C(\bar{x}_{\text{light}})$ in the interest of clarity. b_c is the normalization constant of the chain $c \in \{\text{con}, \text{vis}\}$. To compute a consistent MIS weights throughout a given iteration, we use the normalization constants from the beginning of the iteration. Kelemen et al. [2002] use a similar approach, where they apply MIS to combine samples from a MCMC sampler and independently generated samples (i.e. the large steps).

Restarting MCMC chains. Since the contribution $C(\bar{x}_{\text{light}})$ depends on eye subpaths $\mathcal{E}$, both target functions change when we generate a new set $\mathcal{E}$. To ensure convergence, we restart the MCMC chain at the beginning of each iteration. We then average the results of individual iterations to get the final result. The same algorithm was employed by Hachisuka and Jensen [2011] for photon tracing.

5.3 MCMC Sampler Implementation

Algorithm 2 provides a detailed pseudocode of our MCMC algorithm used to sample light subpaths. It receives the precomputed set of eye subpaths $\mathcal{E}$, and starts by initializing the two Markov chains (line 1). The loop generating individual light subpaths starts by picking (at random) the eye subpath $\bar{x}_{\text{eye}}$ (line 3) that the next light subpath will connect to. After that, we select the Markov chain c to update (line 4): we alternate between the contribution (con) and visibility (vis) chains, both receiving the same number of samples.

To propose a candidate u for the selected chain c, we mutate the current state $u_{c,j-1}$ in the primary sample space (line 5). Once we have the primary sample space candidate u, we map it to a light subpath $\bar{x}_{\text{light}}$ by using a standard photon tracing algorithm that uses u as a vector of random numbers (line 6). The light subpath contribution C is then calculated using Eq. (3) (line 7). We then use the contribution C to evaluate the target function for the proposed light subpath using Eq. (4) or (5), which is, in turn, used to calculate the acceptance probability a (line 8). Based on that probability, we either accept or reject the proposed light subpath (lines 10–12).

We accumulate the contributions of both the proposed and the current light subpath (lines 14–15), as per the use of expected values [Veach and Guibas 1997]. To combine the contributions from the contribution and visibility chains, we apply the MCMC MIS weights defined by Eq. (7). The replica exchange probability

Algorithm 2 Our MCMC algorithm for generating light subpaths.

1: $[u_{\text{con},0}, u_{\text{vis},0}] :=$ INITIALIZECHAINS()
2: **for** $j = 1 \ldots N$ **do**
3: $\quad \bar{x}_{\text{eye}} :=$ GETNEXTEYEPATH($\mathcal{E}$)
4: $\quad$ Select $c \in \{\text{con}, \text{vis}\}$ // alternate between contrib. and visibility chains
5: $\quad u :=$ MUTATE($u_{c,j-1}$) // propose candidate via mutation
6: $\quad \bar{x}_{\text{light}} :=$ TRACELIGHTSUBPATH(u)
7: $\quad C :=$ COMPUTESUBPATHCONTRIB($\bar{x}_{\text{light}}$, $\bar{x}_{\text{eye}}$, $\mathcal{E}$) // Eq. (3)
8: $\quad a :=$ CALCACCEPTANCEPROB($p_c^*(C)$, $p_c^*(C_{c,j-1})$) // Eq. (1)
9: $\quad$ **if** $a >$ RANDOM() **then**
10: $\quad\quad [u_{c,j}, C_{c,j}] := [u, C]$ // accept mutation
11: $\quad$ **else**
12: $\quad\quad [u_{c,j}, C_{c,j}] := [u_{c,j-1}, C_{c,j-1}]$ // reject mutation, keep old state
13: $\quad$ **end if**
14: $\quad \text{image}_c \mathrel{+}= (1-a) \times \frac{C_{c,j-1}}{p_c^*(C_{c,j-1})} \times$ MCMCMISWEIGHT($C_{c,j-1}$)
15: $\quad \text{image}_c \mathrel{+}= a \times \frac{C}{p_c^*(C)} \times$ MCMCMISWEIGHT(C)
16: $\quad r := \min\left\{1, \frac{p^*_{\text{con}}(C_{\text{vis},j})}{p^*_{\text{con}}(C_{\text{con},j})}\right\}$ // exchange probability, Eq. (2)
17: $\quad$ **if** $r >$ RANDOM() **then**
18: $\quad\quad$ SWAP($[u_{\text{vis},j}, C_{\text{vis},j}]$, $[u_{\text{con},j}, C_{\text{con},j}]$) // replica exchange
19: $\quad$ **end if**
20: $\quad$ **if** largeStep $=$ true **then**
21: $\quad\quad$ UPDATENORMALIZATION(C) // Eq. (6)
22: $\quad$ **end if**
23: **end for**
24: OUTPUT($2\frac{(b_{\text{vis}}\text{image}_{\text{vis}} + b_{\text{con}}\text{image}_{\text{con}})}{NM}$) // progressive output

is calculated on line 16 by Eq. (2), simplified by the fact that $p^*_{\text{vis}}(C(\bar{x}_{\text{light}})) = 1$ for any accepted light subpath. The probabilistic state exchange is then carried out on line 18. Finally, whenever an independent, large step, is performed, we update normalization constant estimates $\langle b_{\text{vis}} \rangle$ and $\langle b_{\text{con}} \rangle$, Eq. (6) (line 21). While large steps are only generated for the visibility chain, we can use them to progressively update normalization constants of both chains during the entire rendering process.

After having generated N light subpaths in the current iteration, we combine their contribution and output the progress of rendering (line 24), while dividing by the number of iterations M.

5.4 MIS Weight Computation

We now return to the computation of MIS weights for combining the VCM/UPS path sampling techniques in Eq. (3). A standard approach is to use the balance heuristic [Veach and Guibas 1995]

$$w_t(\bar{x}) = \frac{\hat{p}_t(\bar{x})}{\sum_{t' \in \mathcal{T}} \hat{p}_{t'}(\bar{x})}, \tag{8}$$

where $\mathcal{T}$ is the set of all VCM/UPS sampling techniques that can generate the path $\bar{x}$. $\hat{p}_t(\bar{x})$ denotes the effective pdf with which the full path $\bar{x}$ would be sampled by the path sampling technique t.

Computing exact $\hat{p}_t(\bar{x})$ in our algorithm is unfeasible in practice. This is a shared limitation with PSSMLT, and has not been reported by previous work. The issue lies in the fact that any given full path $\bar{x}$ can be generated by infinitely many different primary samples $\mathcal{U}_{\bar{x}} \subseteq \mathcal{U}$, each of which corresponds to a *light* subpath with the same prefix (see Fig. 10). Furthermore, each such $u \in \mathcal{U}_{\bar{x}}$ usually has a different target function value, since it is given by the sum of the contributions of all full paths generated by the light subpath. The exact pdf $\hat{p}_t(\bar{x})$ is therefore given by

$$\hat{p}_t(\bar{x}) = p_t(\bar{x}) p_{\mathcal{U},c}(\bar{x}) = p_t(\bar{x}) \int_{\mathcal{U}_{\bar{x}}} \frac{p_c^*(C^*(u))}{b_c} d\pi_{\bar{x}}(u). \tag{9}$$

Here $p_t(\bar{x})$ is the pdf of the path sampling technique t, i.e. the Jacobian of mapping u to $\bar{x}$, which is the same for all $u \in \mathcal{U}_{\bar{x}}$. $p_{\mathcal{U},c}(\bar{x})$ is the marginalized probability density of sampling $u \in \mathcal{U}_{\bar{x}}$ by the Markov chain c. $p_c^*(C^*(u))$ is the target function of the chain c, and b_c is the normalization of the chain c. The measure function $\pi_{\bar{x}}(u)$ is then a simple volume measure which ignores the common prefix of all $u \in \mathcal{U}_{\bar{x}}$.

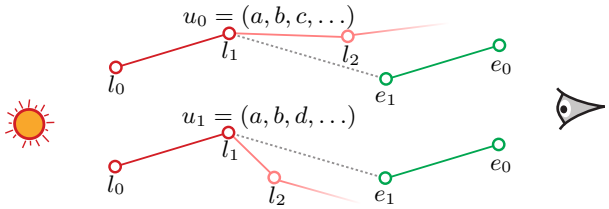

Figure 10: *A single full path $\bar{x} = (l_0, l_1, e_1, e_0)$ can be generated by infinitely many different primary samples (u_0 and u_1 in this example) that all share the same prefix (a, b). The remaining components of each primary sample ($c, \ldots$ and $d, \ldots$ in this example) are used to generate the vertices of the light subpath after l_1. Those are* not *part of the full path, but still affect its pdf.*

PSSMLT approximates the pdf in Eq. (9) by ignoring $p_{\mathcal{U},c}(\bar{x})$ (i.e. $\hat{p}_t(\bar{x}) \approx p_t(\bar{x})$). While such an approach is suboptimal, because the MIS weights do not use the actual pdf of the paths, it usually works well in practice. For this reason, we adopt it for paths generated by our *contribution chain* as well.

However, we show that for the *visibility chain*, we can compute the exact path pdf and thus improve the algorithm efficiency (see Fig. 11). While a full path $\bar{x}$ can be constructed from infinitely many primary samples $\mathcal{U}_{\bar{x}}$ of the visibility chain, they all have the same target function value equal to one. Given the above and the fact that $\mathcal{U}_{\bar{x}}$ has a unit volume, the exact pdf $\hat{p}_t(\bar{x})$ is computed as

$$\hat{p}_t(\bar{x}) = p_t(\bar{x}) \int_{\mathcal{U}_{\bar{x}}} \frac{d\pi_{\bar{x}}(u)}{b_{\text{vis}}} = \frac{p_t(\bar{x})}{b_{\text{vis}}}. \tag{10}$$

Please note that the paths generated by the path tracing techniques are independent of the MCMC process, and thus their pdfs are computed as in bidirectional path tracing (i.e. $\hat{p}_t(\bar{x}) = p_t(\bar{x})$).

6 Results

We have implemented our algorithm and baseline VCM/UPS on top of the Mitsuba renderer [Jakob 2010]; the code will be released upon publication. We ran all the experiments on a PC with an Intel Core i7 at 3.50 GHz with 16 GB RAM using eight logical cores. The reference images were rendered using the VCM/UPS algorithm over the course of two weeks.

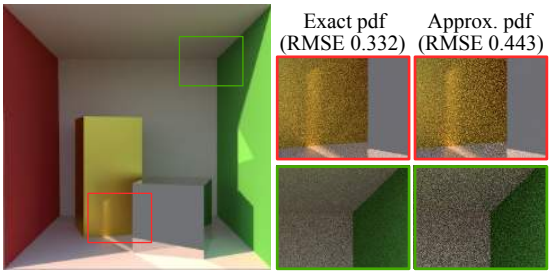

Figure 11: *Equal-sample (64 iterations) comparison of results generated using only the visibility chain with the exact pdfs used for MIS weights computation (left) and with the approximated pdfs (right). Using exact pdfs for MIS weights computation decreases noise in the image due to better weighting of sampling techniques.*

Figure 12: *Equal-time comparisons (1 hour) of our method, baseline VCM/UPS, and MLT with manifold exploration (MEMLT). The reported speedup is calculated as the squared ratio of RMSE values, and gives the expected relative time that the baseline VCM/UPS would take to attain the RMSE of each respective algorithm. EV is the exposure value adjustment of the insets over the base image. The main text discusses the important features of the results shown here.*

From among the previous MCMC algorithms we have tested, MLT with manifold exploration (MEMLT) [Jakob and Marschner 2012] provided the best results which is why we use it for most of our comparisons in the paper. However, additional results not included in the paper can be found in the supplemental material. MEMLT parameters used in Mitsuba have been manually tuned for each scene to ensure the best performance. We used adaptive mutation size [Hachisuka and Jensen 2011] for our algorithm, PSSMLT, and Multiplexed MLT (MMLT) [Hachisuka et al. 2014] with the target acceptance rate of 23.4%, which is roughly attained in our tests.

Radii for the vertex merging (i.e. photon density estimation) techniques at all the vertices of an eye subpath are determined by the pixel footprint at the first vertex of the respective subpath. We have not used any radius reduction [Hachisuka et al. 2008] but we have experimentally verified that the radius is small enough so that any bias is numerically and visually negligible in our scenes. Each iteration of our algorithm and baseline VCM/UPS uses one eye and light subpath per pixel.

Comparison to VCM/UPS and MLT. Fig. 1 and Fig. 12 show equal-time (one hour) comparisons of our algorithm, baseline VCM/UPS, and MEMLT. All three cases represent a typical scene configuration where an interior is illuminated from the outside via an opening such as a window. Table 1 summarizes rendering parameters and statistics. The average replica exchange success rate of about 27% roughly corresponds to the optimal values reported in the literature [Atchadé et al. 2010]; this provides an empirical evidence of the effectiveness of our two-chain design.

	Kitchen	Living room	Classroom
Image resolution	1024 × 768	1024 × 768	960 × 480
#iterations VCM/UPS	675	641	1002
#iterations our algorithm	357	406	525
avg. exchange rate (our algo.)	26.5%	26.4%	27.6%

Table 1: *Settings and statistics for scenes in Fig. 1 and 12.*

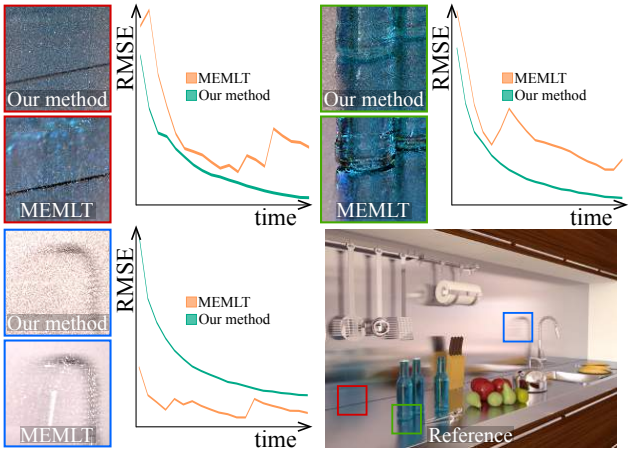

Figure 13: *RMS error plots of our method and MEMLT calculated for the insets of the Kitchen scene in Fig. 1. While MEMLT may have lower numerical error than our method in some cases, its convergence behavior can be highly unpredictable. Our method predictably converges toward a noise-free solution.*

Overall, the performance of baseline VCM/UPS suffers from an insufficient number of light subpaths due to complex visibility. Even for the **Living room** scene which features a relatively large window and is thus easier to render with baseline VCM/UPS, our algorithm achieves a speedup of more than $4\times$. MEMLT tends to fail to evenly explore the path space due to the presence of highly glossy/specular reflections. For example, the insets of the **Living room** scene highlight how MEMLT can either completely miss or over-explore specular-diffuse-specular paths. Our algorithm consistently outperforms both of those state-of-the art algorithms in these comparisons.

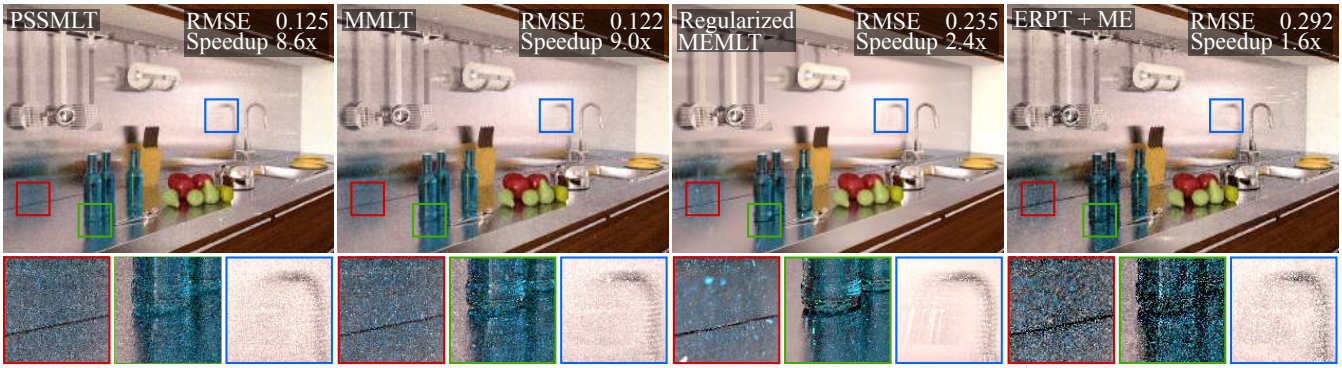

Figure 14: *An equal-time (one hour) comparison of state-of-the-art MCMC-based algorithms in the Kitchen scene. The algorithms shown are PSSMLT [Kelemen et al. 2002], Multiplexed MLT (MMLT) [Hachisuka et al. 2014], MEMLT with additional regularization [Kaplanyan and Dachsbacher 2013], and energy redistribution path tracing with manifold exploration (ERPT + ME) [Cline et al. 2005]. The speedup is calculated as the squared ratio of RMSE values, and gives the expected relative time that the baseline VCM/UPS would take to attain the RMSE of each respective algorithm. Refer to Fig. 1 for the results of baseline VCM/UPS and our algorithm.*

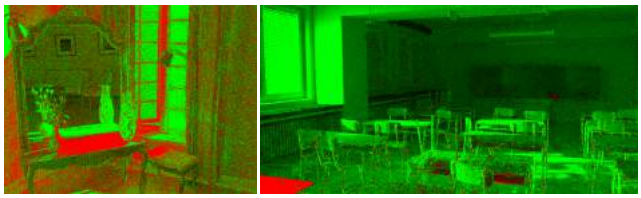

Figure 15: *Positive (red)-negative (green) difference of the MEMLT results in Fig. 12 from the respective references for the Living room (left) and Classroom (right) scenes.*

The **Kitchen** scene in Fig. 1 features a counter-top of highly anisotropic metal which often leads to various image artifacts due to sample correlation in MEMLT. Fig. 13 shows RMSE plots within the insets of this scene, highlighting MEMLT's uneven convergence. Our algorithm steadily converges toward the correct solution within all the insets. We found that our algorithm generally exhibits this predictable convergence behavior.

While the images rendered by MEMLT tend to be less noisy than from other algorithms, they are not necessarily converged to the correct solution. Fig. 15 highlights this problem in the **Living room** and **Clasroom** scenes. While the noise-free image of MEMLT in Fig. 12 could be mistakenly considered converged, its RMS error as well as the visualization in Fig. 15 show that the contrary is true.

Comparison to other state-of-the art algorithms. Fig. 14 shows a comparison of our method to other state-of-the art MCMC-based algorithms in the Kitchen scene (similar results for other scenes are included in the supplemental material). Namely, we compare to PSSMLT [Kelemen et al. 2002], Multiplexed MLT (MMLT) [Hachisuka et al. 2014], MEMLT with additional regularization [Kaplanyan and Dachsbacher 2013], energy redistribution path tracing with manifold exploration (ERPT + ME) [Cline et al. 2005]. For ERPT + ME we set the average number of chains per pixel to one and limit the number of mutations per chain to 100 (default setting in Mitsuba). ERPT + ME strives to improve stratification compared to MEMLT, but its short Markov chains impair its efficiency. For the regularized MEMLT, we use our own implementation, which performs regularization of interactions on highly glossy materials by adjusting their roughness (this idea was suggested in the original paper). While the strength of regularization differs for each material, the shrinkage parameter is set globally to 10^{-4}. The regularized MEMLT allows easier sampling of some types of paths, however it does not improve MEMLT's irregular convergence. In fact, the results show that all those algorithms suffer from severe image artifacts due to sample correlation in MCMC.

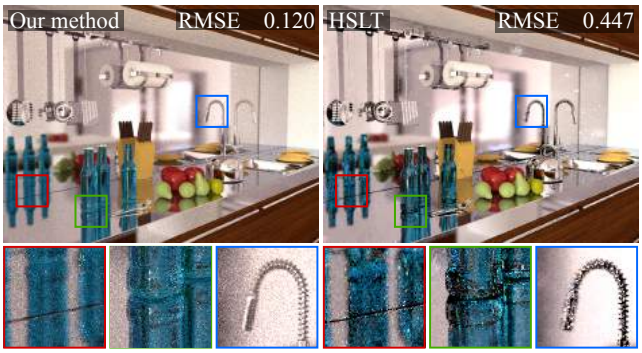

Figure 16: *An equal-time (one hour) comparison of our method and half vector space light transport (HSLT) [Hanika et al. 2015] in the Kitchen scene with isotropic materials. We use this comparison because the HSLT implementation in Mitsuba cannot handle anisotropic materials. Reference is in the supplemental material.*

Fig. 16 shows a comparison to half vector space light transport (HSLT) [Hanika et al. 2015] in a variant of the Kitchen scene. Since the publicly available implementation of HSLT (in Mitsuba) cannot handle anisotropic materials, we set all materials to be isotropic for the sake of this specific comparison. We have set the universal perturbation probability in HSLT to get an overall 23.4% acceptance ratio. While previous work has reported better results for HSLT than MEMLT, we have not been able to observe any such improvement in our tests (even though we have consulted HSLT settings with the authors). In this specific example, we see that HSLT shares the issues of other MCMC methods: it suffers from irregular convergence and fails to explore glossy/specular light transport.

Temporal coherence. To compare temporal coherence of PSSMLT, MEMLT, VCM/UPS and our algorithm, we have generated videos (using equal time renderings) showing a short flythrough of the Kitchen scene (the videos are in the supplemental material). While PSSMLT and MEMLT exhibit pronounced flickering artifacts, our results are substantially more well-behaved, only showing the high-frequency noise typical for regular Monte Carlo rendering. The results of VCM/UPS maintain good temporal coherence, however the overall noise level is very high.

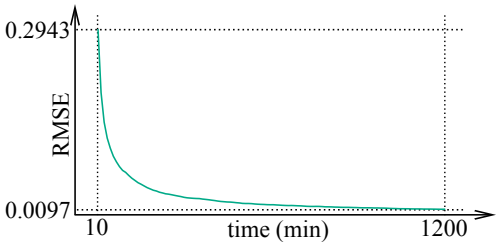

Figure 17: *Empirical verification of our algorithm's convergence in the Kitchen scene. The figure shows a steadily decreasing plot of RMSE during 20 hours of rendering.*

Consistency and behavior in simple scenes. Fig. 17 plots the RMSE of our algorithm over 20 hours of rendering in the Kitchen scene. The RMSE steadily approaches to zero, which provides an empirical evidence of the convergence of our method.

We have mentioned earlier and shown in Fig. 4 the low overhead of our algorithm over baseline VCM/UPS in simple scenes. Fig. 18 additionally demonstrates this in a comparison with PSSMLT. Simpler algorithms like path tracing would outperform both our algorithm and VCM/UPS in a very simple scenes, however it is at the cost of their lower robustness.

Convergence analysis of our algorithm. In order to analyze which parts of our algorithm are responsible for its robustness and uniform convergence, we have run our algorithm with three different features switched on or off. Fig. 19 describes the different versions and shows their RMSE plots (the images rendered by the different versions are then shown in Appendix A). We can observe that some of the plots are not monotonic, which is caused by the algorithm getting stuck in local maxima and thus temporarily diverging from the reference. Clearly, the versions of our algorithm that combine two chains via replica exchange (green lines) are more robust to this problem. The plots also show that while using stratified eye subpaths lowers RMSE compared to using eye subpaths generated by MCMC, it does not guarantee monotonicity of the convergence plot. Using all VCM/UPS techniques instead of just bidirectional path tracing techniques then in most cases results in lower RMSE. We therefore conclude that the uniform convergence of our algorithm is mainly due to the use of replica exchange.

7 Limitations and Future Work

Though our algorithm robustly handles various scene configurations that are inefficient to render with existing algorithms, there are several limitations that need to be resolved by further work.

Negative effects of independent eye subpaths. While using stratified, independently generated eye subpaths is one of the key

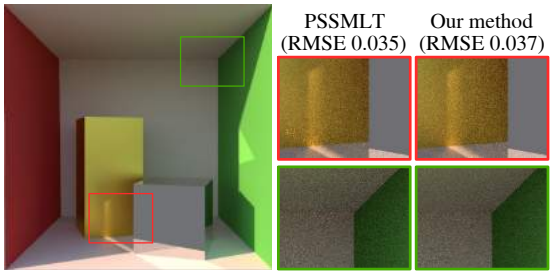

Figure 18: *Our method has negligible overhead over PSSMLT in this simple scene, where the latter algorithm works well. The images took 2 minutes to render.*

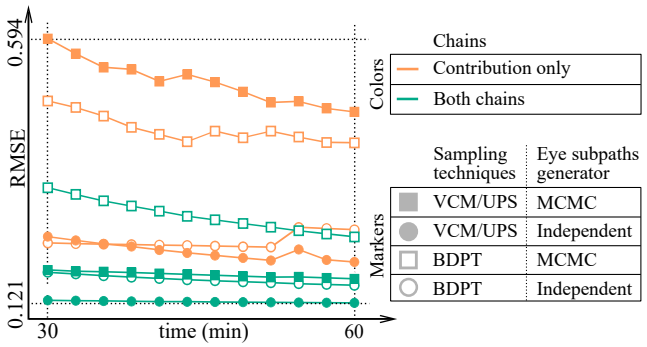

Figure 19: *Here we show convergence plots for our algorithm in the Kitchen scene with three different features switched on or off. First, we can run our algorithm only with the contribution chain (orange lines) or with both chains using replica exchange (green lines). Second, we can either enable all VCM/UPS techniques (full markers) or enable only bidirectional path tracing (BDPT) techniques (empty markers). Finally, we can generate eye subpaths using a MCMC algorithm (square markers) or independent Monte Carlo (circle markers). We can observe from the plots that all components of our algorithm (namely replica exchange, all VCM/UPS techniques, independent eye subpaths) contribute to achieving the low RMSE. The plots are further discussed in the text.*

ideas behind the efficiency of our algorithm, it can negatively affect local exploration in our MCMC sampler. Consider a light subpath that makes a contribution only via a single bidirectional connection to the currently selected eye subpath. Even when the MCMC sampler proposes only a slight mutation of the *light* subpath, our algorithm still takes a completely independent *eye* subpath to connect to. As such, the previously successful bidirectional connection is likely to become invalid. Fortunately, this negative effect is diminished by the presence of vertex merging techniques, which do not depend on the currently selected eye path at all (see Fig. 20).

Sub-optimal MIS weights. Our MIS weight calculation takes an approximated pdf (as in PSSMLT) for paths generated by the contribution chain (see Sec. 5.4). Though the practical impact of this approximation is unknown, since it is unfeasible to compute the exact pdf, deriving new MIS weights based on true pdfs with which our algorithm samples the paths would certainly increase its efficiency. Furthermore, we currently base our MIS weights on balance heuristic, which has been proven to work well under the condition of independent samples. This condition is however broken for MCMC and therefore using balance heuristic might be sub-optimal.

Near-specular paths. One limitation of our algorithm, shared by other approaches, is handling nearly specular paths. That is, paths

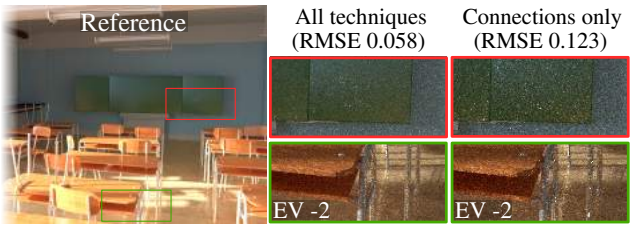

Figure 20: *Equal-samples comparison (525 iterations) of our method using all VCM/UPS techniques and using only bidirectional techniques (connections). Including vertex merging techniques diminishes the negative effects of using independent eye subpaths in the MCMC sampler.*

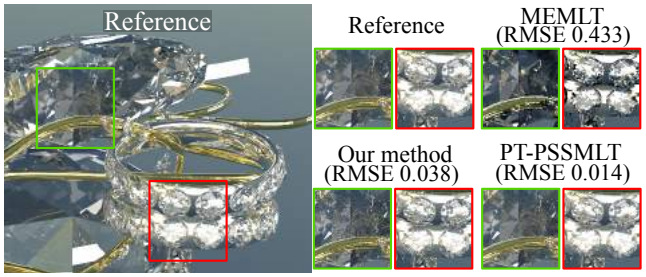

Figure 21: *Equal-time comparison (10 minutes) of rendering a scene with only purely specular materials. While our algorithm delivers worse results than PSSMLT over unidirectional path-tracing (PT-PSSMLT), it outperforms MEMLT, which gets stuck in local maxima of its target function.*

that interact *only* with highly specular materials. Since such paths can be effectively sampled only by unidirectional path tracing technique, our algorithm provides no advantage over existing solutions in handling them (see Fig. 21).

Future Work. In our algorithm, we map each primary sample to a set of full paths using all the VCM/UPS path sampling techniques. Hachisuka et al. [2014] have shown that using only one selected technique per primary sample can be more efficient. Using a similar approach in our algorithm could potentially increase its efficiency as well. Another interesting topic of future work would be robust rendering of participating media. For instance, our approach could serve as a guideline for applying MCMC sampling to the work of Křivánek et al. [2014a]. Furthermore, claiming that our work completely resolves the problem of unpredictable convergence behavior of MCMC would be an overstatement and there is certainly more work to be done on this topic.

8 Conclusion

We proposed a light transport simulation algorithm which fuses the concepts behind VCM/UPS and MCMC into a practical solution for the first time. Our algorithm can efficiently handle complex glossy and specular transport and complex visibility by taking advantage of both VCM/UPS and MCMC. The design of our algorithm is not a mere combination of the two concepts in a trivial manner. We instead made a number of careful design decisions based on practical insights on the building blocks of the algorithm. More specifically, we designed our algorithm to retain the benefits of combined bidirectional estimators, extensive subpath reuse, and spatial relaxation from VCM/UPS while minimizing the correlation problems of MCMC. The numerical experiments demonstrate that our algorithm indeed features the advantages of both VCM/UPS and MCMC, while suppressing their disadvantages in various scene configurations. Since our algorithm has low overhead over VCM/UPS and primary sample space MLT, it can practically substitute those algorithms in many applications and improve the robustness of light transport simulation.

9 Acknowledgments

We would like to thank the reviewers for their insightful comments, and Johannes Hanika and Anton S. Kaplanyan for their help with setting up half vector space light transport algorithm in Mitsuba. The work was supported by Charles University in Prague, project GA UK 164815, by the grant SVV–2016–260332, by the Czech Science Foundation grant 16–18964S, and by the JSPS KAKENHI Grant Number 15H05308.

References

ATCHADÉ, Y. F., ROBERTS, G. O., AND ROSENTHAL, J. S. 2010. Towards optimal scaling of Metropolis-coupled Markov chain Monte Carlo. *Statistics and Computing 21*, 4, 555–568.

BASHFORD-ROGERS, T., DEBATTISTA, K., AND CHALMERS, A. 2012. A significance cache for accelerating global illumination. *Computer Graphics Forum 31*, 6, 1837–51.

BROOKS, S., GELMAN, A., JONES, G. L., AND MENG, X.-L. 2011. *Handbook of Markov Chain Monte Carlo*. Chapman and Hall/CRC.

CHEN, J., WANG, B., AND YONG, J.-H. 2011. Improved stochastic progressive photon mapping with metropolis sampling. *Comput. Graph. Forum (EGSR 2011) 30*, 4.

CLINE, D., TALBOT, J., AND EGBERT, P. 2005. Energy redistribution path tracing. *ACM Trans. Graph. (SIGGRAPH 2005) 24*, 3.

DACHSBACHER, C., KŘIVÁNEK, J., HAŠAN, M., ARBREE, A., WALTER, B., AND NOVÁK, J. 2014. Scalable realistic rendering with many-light methods. *Comput. Graph. Forum 33*, 1.

FAN, S., CHENNEY, S., AND LAI, Y.-C. 2005. Metropolis photon sampling with optional user guidance. In *Eurographics Symposium on Rendering (EGSR '05)*, 127–138.

GEORGIEV, I., KŘIVÁNEK, J., DAVIDOVIČ, T., AND SLUSALLEK, P. 2012. Light transport simulation with vertex connection and merging. *ACM Trans. Graph. (SIGGRAPH Asia '12) 31*, 6.

HACHISUKA, T., AND JENSEN, H. W. 2009. Stochastic progressive photon mapping. *ACM Trans. Graph. (SIGGRAPH Asia 2009) 28*, 5.

HACHISUKA, T., AND JENSEN, H. W. 2011. Robust adaptive photon tracing using photon path visibility. *ACM Trans. Graph. 30*, 5.

HACHISUKA, T., OGAKI, S., AND JENSEN, H. W. 2008. Progressive photon mapping. *ACM Trans. Graph. (SIGGRAPH Asia '08) 27*, 5.

HACHISUKA, T., PANTALEONI, J., AND JENSEN, H. W. 2012. A path space extension for robust light transport simulation. *ACM Trans. Graph. (SIGGRAPH Asia '12) 31*, 6.

HACHISUKA, T., KAPLANYAN, A. S., AND DACHSBACHER, C. 2014. Multiplexed Metropolis light transport. *ACM Trans. Graph. (SIGGRAPH 2014) 33*, 4 (July), 100:1–100:10.

HANIKA, J., KAPLANYAN, A., AND DACHSBACHER, C. 2015. Improved half vector space light transport. *Comput. Graph. Forum (EGSR 2015) 34*, 4.

HASTINGS, W. K. 1970. Monte Carlo sampling methods using Markov chains and their applications. *Biometrika 57*, 1, 97–109.

HOBEROCK, J., AND HART, J. C. 2010. Arbitrary importance functions for Metropolis light transport. *Comput. Graph. Forum 29*, 6, 1993–2003.

JAKOB, W., AND MARSCHNER, S. 2012. Manifold exploration: A Markov chain Monte Carlo technique for rendering scenes with difficult specular transport. *ACM Trans. Graph. 31*, 4.

JAKOB, W., 2010. Mitsuba renderer. http://mitsuba-renderer.org.

JAROSZ, W., NOWROUZEZAHRAI, D., THOMAS, R., SLOAN, P.-P., AND ZWICKER, M. 2011. Progressive photon beams. *ACM Trans. Graph. (SIGGRAPH Asia 2011) 30*, 6.

JENSEN, H. W. 1995. Importance driven path tracing using the photon map. In *Eurographics Workshop Rendering*, 326–335.

JENSEN, H. W. 1996. Global illumination using photon maps. In *Eurographics Workshop on Rendering*, Springer-Verlag, 21–30.

KAPLANYAN, A. S., AND DACHSBACHER, C. 2013. Path space regularization for holistic and robust light transport. *Comput. Graph. Forum (Eurographics 2013) 32*, 2, 63–72.

KELEMEN, C., SZIRMAY-KALOS, L., ANTAL, G., AND CSONKA, F. 2002. A simple and robust mutation strategy for the Metropolis light transport algorithm. *Comp. Graph. Forum (Eurographics 2002) 21*, 3, 531–540.

KITAOKA, S., KITAMURA, Y., AND KISHINO, F. 2009. Replica exchange light transport. *Computer Graphics Forum 28*, 8.

KNAUS, C., AND ZWICKER, M. 2011. Progressive photon mapping: A probabilistic approach. *ACM Trans. Graph. 30*, 3.

KŘIVÁNEK, J., GEORGIEV, I., HACHISUKA, T., VÉVODA, P., ŠIK, M., NOWROUZEZAHRAI, D., AND JAROSZ, W. 2014. Unifying points, beams, and paths in volumetric light transport simulation. *ACM Trans. Graph. (SIGGRAPH 2014) 33*, 4.

KŘIVÁNEK, J., KELLER, A., GEORGIEV, I., KAPLANYAN, A., FAJARDO, M., MEYER, M., NAHMIAS, J.-D., KARLÍK, O., AND CANADA, J. 2014. Recent advances in light transport simulation: Some theory and a lot of practice. In *ACM SIGGRAPH 2014 Courses*.

LAFORTUNE, E. P., AND WILLEMS, Y. D. 1993. Bi-directional path tracing. In *Proc. of Compugraphics 93*.

LI, T.-M., LEHTINEN, J., RAMAMOORTHI, R., JAKOB, W., AND DURAND, F. 2015. Anisotropic Gaussian mutations for Metropolis light transport through Hessian-Hamiltonian dynamics. *ACM Trans. Graph. (SIGGRAPH Asia 2015) 34*, 6.

PAULY, M., KOLLIG, T., AND KELLER, A. 2000. Metropolis light transport for participating media. In *Proceedings of the Eurographics Workshop on Rendering Techniques 2000*, 11–22.

POPOV, S., RAMAMOORTHI, R., DURAND, F., AND DRETTAKIS, G. 2015. Probabilistic connections for bidirectional path tracing. *Computer Graphics Forum (Proc. of EGSR) 34*, 4.

SEGOVIA, B., IEHL, J.-C., AND PÉROCHE, B. 2007. Metropolis instant radiosity. *Comput. Graph. Forum (Eurographics 2007) 26*, 3, 425–434.

SWENDSEN, R. H., AND WANG, J.-S. 1986. Replica Monte Carlo simulation of spin-glasses. *Phys. Rev. Lett. 57*, 2607–2609.

VEACH, E., AND GUIBAS, L. 1994. Bidirectional estimators for light transport. In *Proc. Eurographics Rendering Workshop*.

VEACH, E., AND GUIBAS, L. J. 1995. Optimally combining sampling techniques for Monte Carlo rendering. In *SIGGRAPH '95*.

VEACH, E., AND GUIBAS, L. J. 1997. Metropolis light transport. In *SIGGRAPH '97*.

VEACH, E. 1997. *Robust Monte Carlo methods for light transport simulation*. PhD thesis, Stanford University.

VORBA, J., KARLÍK, O., ŠIK, M., RITSCHEL, T., AND KŘIVÁNEK, J. 2014. On-line learning of parametric mixture models for light transport simulation. *ACM Trans. Graph. (SIGGRAPH '14) 33*, 4.

ZWICKER, M., JAROSZ, W., LEHTINEN, J., MOON, B., RAMAMOORTHI, R., ROUSSELLE, F., SEN, P., SOLER, C., AND YOON, S.-E. 2015. Recent advances in adaptive sampling and reconstruction for Monte Carlo rendering. *Computer Graphics Forum (Proceedings of Eurographics) 34*, 2 (May), 667–681.

A Algorithm components comparison

Fig. 22 shows images rendered by the different versions of our algorithm (see Fig. 19 for the description of the different versions).

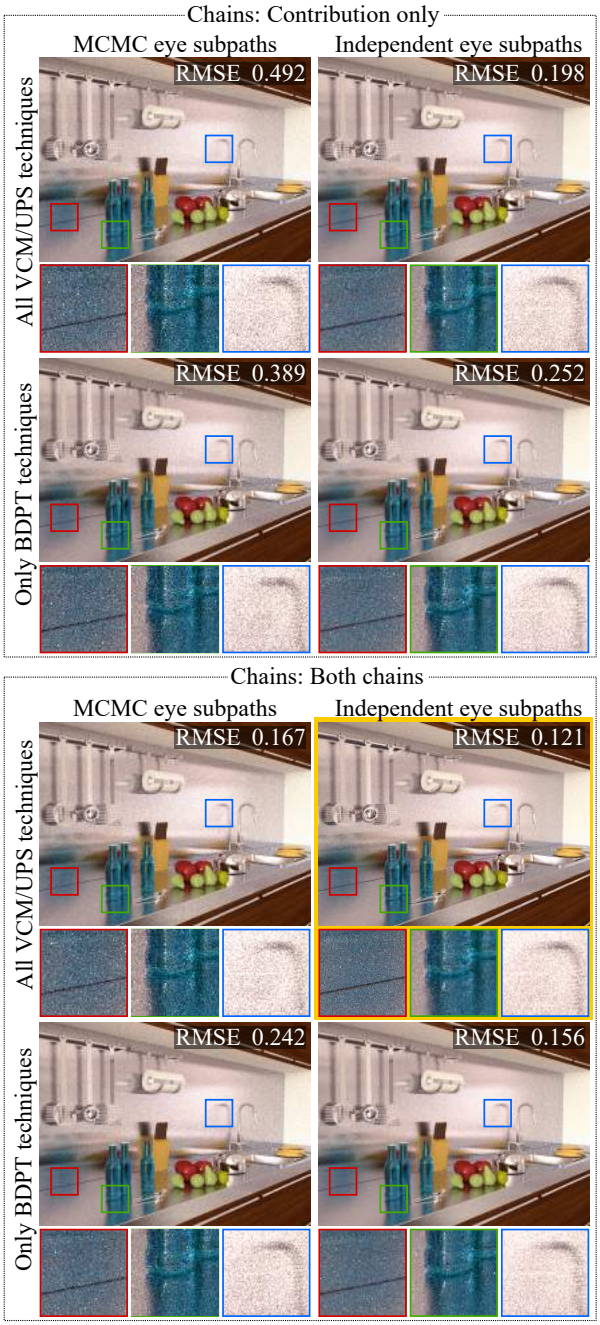

Figure 22: *An equal-time (one hour) comparison of different variants of our algorithm. The final most robust method is marked by the orange rectangle.*