Improving Global Exploration of MCMC Light Transport Simulation

Martin Šik* Jaroslav Křivánek†
Charles University in Prague

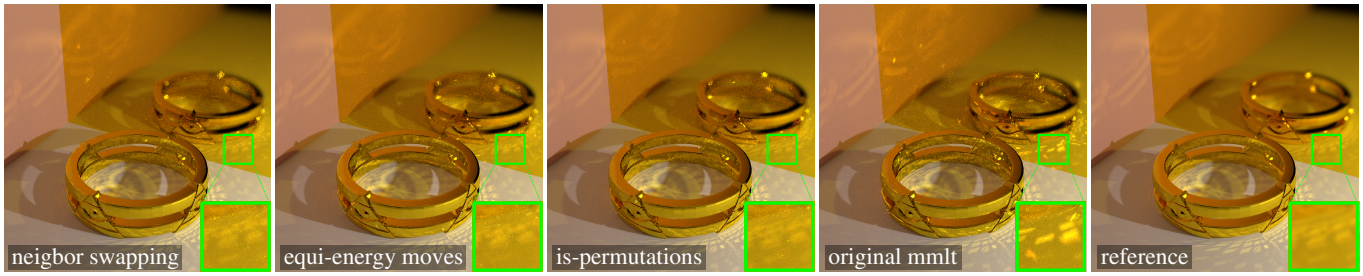

Figure 1: *Here we show an equal-time comparison of the original MMLT (d) to three variants of our algorithm – Neighbor swapping (a), Equi-energy moves (b), Importance-sampled permutations (c). Reference is shown in the image (e). Due to insufficient global exploration in original MMLT some of the transport (e.g. reflected caustics) is missing in the image (d), while our algorithm is able to recover it.*

1 Introduction

Markov Chain Monte Carlo (MCMC) has recently received a lot of attention in light transport simulation research [Hanika et al. 2015; Hachisuka et al. 2014]. While these methods aim at high quality sampling of local extremes of the path space (so called *local exploration*), the other issue – discovering these extremes – has been so far neglected. Poor *global exploration* results in oversampling some parts of the paths space, while undersampling or completely missing other parts (see Fig. 1). Such behavior of MCMC-based light transport algorithms limits their use in practice, since we can never tell for sure whether the image has already converged.

Outside of computer graphics, the problem of global exploration has received much attention. One of the most popular methods for improving global exploration of MCMC is *parallel tempering* (PT) [Swendsen and Wang 1986]. While PT has already been applied in light transport simulation [Kitaoka et al. 2009; Hachisuka and Jensen 2011], we believe its full potential is yet to be exploited. In our work we have developed a parallel tempering algorithm specifically suited for light transport simulation. We demonstrate its three variants and show that they improve global exploration of MCMC algorithms in light transport simulation.

Keywords: rendering, global illumination, MCMC

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

2 Background and Previous Work

A standard MCMC algorithm, Metropolis-Hastings (MH) [Hastings 1970], generates samples from a given target function by mutating a current sample and then accepting a new one with a given probability. However, when a target function has many local maxima and/or discontinuities, the algorithm tends to get stuck in a small region of the sampled domain and other regions are undersampled or not discovered at all. Parallel tempering improves MH algorithm by running several parallel chains. In the simplest case, we have a main chain with the original target function and a secondary chain that uses its flattened version. The secondary chain more easily explores the sampled domain and helps the main chain by swapping their current samples with a given probability. To improve the algorithm efficiency we can use more secondary chains with increasingly flattened target functions.

MH was introduced to light transport simulation as Metropolis Light Transport (MLT) [Veach and Guibas 1997]. We however base our algorithm on Multiplexed MLT (MMLT) [Hachisuka et al. 2014]. MMLT mutates samples in space of random numbers (Primary sample space [Kelemen et al. 2002]), which are then mapped using a selected bidirectional path sampling technique to the path space. Thanks to this mapping, the target function is more flat and thus easier to explore than the one used in MLT.

Parallel tempering (a.k.a. replica-exchange) was introduced to light transport simulation by Kitaoka et al [2009]. In their method, each chain is used to efficiently explore a different region of the sampled domain. However, many regions remain complicated for all of the chains and thus the efficiency of PT is greatly reduced. In our algorithm, we flatten the target function in all regions.

3 Tempering of Light Transport

The first step in developing the PT algorithm is to decide how to flatten/simplify the target function of the secondary chains. In many physics simulations the complexity of the target function is dependent on temperature, increasing temperature (tempering) of the simulation leads to flatter target functions. In light transport simulation we can observe the same behavior with glossiness/roughness of the used bidirectional reflectance functions (BRDF). Increasing roughness of BRDF leads to regularization of the simulation [Kaplanyan and Dachsbacher 2013] and to flatter target functions.

Unfortunately, increasing roughness for all surfaces by even a small amount leads to a target function that differs too much from the original, which in turn results in low efficiency of PT. To avoid this issue, we only roughen BRDF where it matters the most. In MMLT each path is constructed from two subpaths by a connection

*e-mail:sik@cgg.mff.cuni.cz
†e-mail:jaroslav.krivanek@mff.cuni.cz

SIGGRAPH '16, July 24-28, 2016, Anaheim, CA,
ISBN: 978-1-4503-4371-8/16/07
DOI: http://dx.doi.org/10.1145/2945078.2945128

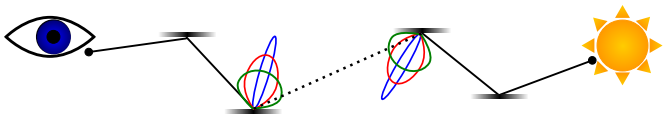

Figure 2: *A full path is constructed from two subpaths (one from the camera, one from the light source) by a connection (dashed line). The main chain uses original BRDF which has a sharp lobe (blue) and thanks to that the connection transfers almost no radiance. The secondary chains use rougher BRDF (red, green) and therefore the amount of transferred radiance is increased and the path is more likely to be accepted.*

between their last vertices. Acceptance of the MCMC mutation mostly depends on the amount of radiance transferred through this connection. We therefore increase roughness of BRDF at the connecting vertices and thus raise the amount of transferred radiance and the chance of accepting the path (see Fig. 2).

4 Chains Swapping

The second crucial aspect of any PT algorithm is the strategy used for swapping of the current samples among chains. For PT to be effective, we must ensure high swap probability. We have developed different strategies to achieve this goal.

- **Neighbor swapping.** The most simple and yet effective strategy is to swap two neighboring chains (i.e. chains with the most similar target functions). High probability of the swap is ensured thanks to their similar target functions.
- **Equi-energy moves.** We borrow this idea, that has previously not been applied in light transport simulation, from a recent statistics paper [Baragatti et al. 2012]. First the codomain of all target functions is split to bins, the so called energy levels. Two chains are then swapped if their target function values lie in the same energy level. This strategy reduces the number of attempted swaps, but on the other hand their probability is very high.
- **Importance-sampled permutations.** Finally, we describe our novel swapping scheme. Unlike the previous techniques, we do not swap only two chains at once; we allow a general permutation of all chains. In order to use the best possible permutation, we select it by importance sampling in such a way that the permutation is always accepted. However, there is a non-zero probability of importance sampling an identity permutation.

5 Implementation and results

We have implemented our algorithm in the Mitsuba renderer [Jakob 2010]. In our experiments we have found that using 5[1] chains is sufficient and we set the roughness to get approximately 23.4% swap probability with the neighbor swapping strategy (we use the same roughness values for the other strategies). Finally, in our scenes we use a microfacet BRDF with the GGX distribution [Walter et al. 2007], which has an intuitive roughness parameter.

We compare our algorithm to the original MMLT algorithm in Fig. 1. The results show that our algorithm with any of the swapping strategies improves global exploration and thus outperforms the original MMLT. Overall we have found the results of neighbor swapping and importance-sampled permutations to be similar. However, they both surpass equi-energy moves due to them having a higher number of successful swaps. More results are shown in the supplemental material.

[1]The original MMLT already uses multiple chains, where each chain generates paths of a fixed path length. In our implementation we use 5 chains for each path length.

6 Conclusion and future work

We have developed an efficient PT algorithm suited for light transport simulation. We have demonstrated three different swapping strategies, two of which have never before been applied in light transport simulation. In our future work, we would like to concentrate on another factor that causes local maxima and discontinues in the target function – scene geometry. While tempering scene geometry is a challenging task, we expect that it will lead to even more efficient MCMC algorithm.

Acknowledgments The work was supported by Charles University in Prague, project GA UK 164815, by the grant SVV–2016–260332, and by the Czech Science Foundation grant 16–18964S.

References

BARAGATTI, M., GRIMAUD, A., AND POMMERET, D. 2012. Parallel tempering with equi-energy moves. *Statistics and Computing 23*, 3, 323–339.

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

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.

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

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.

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. J. 1997. Metropolis light transport. In *SIGGRAPH '97*.

WALTER, B., MARSCHNER, S. R., LI, H., AND TORRANCE, K. E. 2007. Microfacet models for refraction through rough surfaces. In *Proceedings of the 18th Eurographics Conference on Rendering Techniques*, Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, EGSR'07, 195–206.