Computer Graphics Charles University

Improving Global Exploration of MCMC Light Transport Simulation

Martin Šik, Jaroslav Křivánek

Charles University in Prague

SIGGRAPH2016 Render the Possibilities

Problem

Current Markov Chain Monte Carlo (MCMC) methods in light transport simulation aim at high quality sampling of local extremes of the path space, while the other issue – discovering these extremes – has been so far neglected. This results in oversampling some parts of the path space, while undersampling or completely missing other parts. An example of such insufficient **global exploration** is shown below.

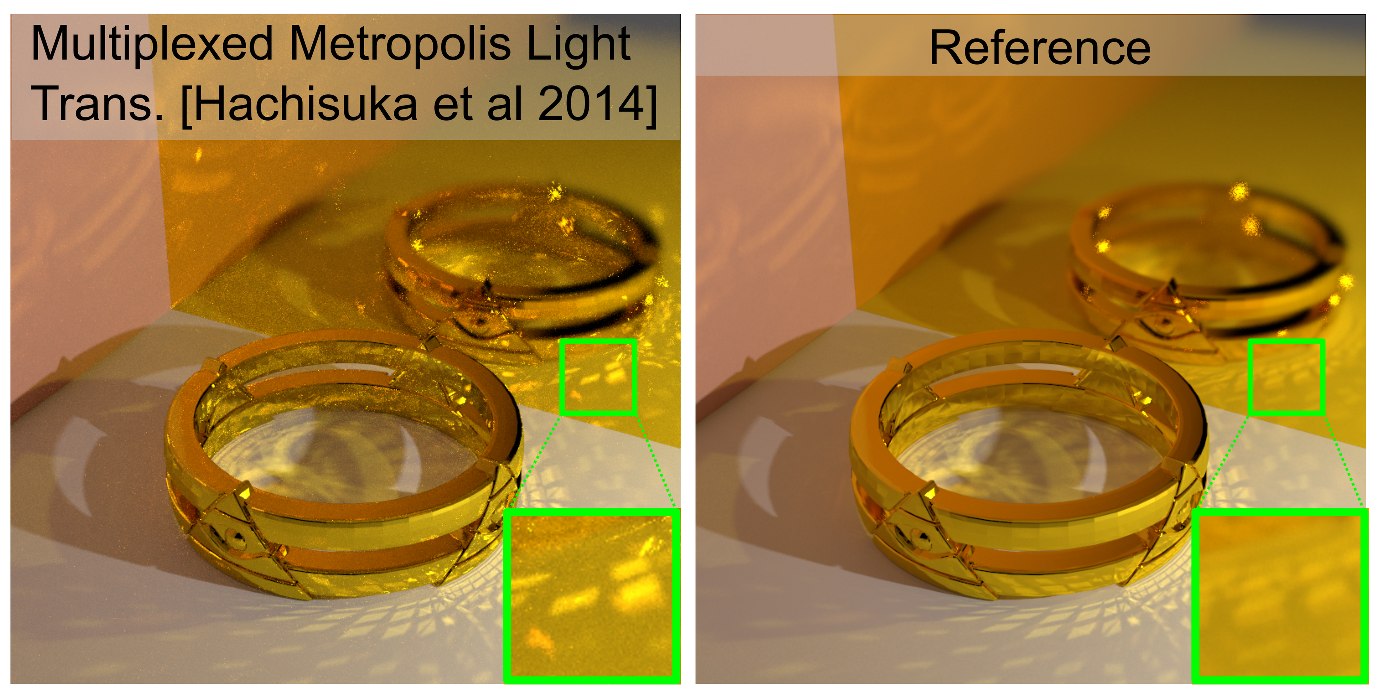

Proposed solution

We overcome the above problem by applying **parallel tempering** [Swendsen and Wang 1986]. Parallel tempering is a general algorithm scheme that improves MCMC algorithms by adding several auxiliary Markov chains with different target functions. The chains' samples are probabilistically swapped to increase global exploration of all chains including the main one.

In our work we design a parallel tempering algorithm specifically suited for light transport and we further investigate three different chain swapping strategies.

Related work

Parallel tempering has been previously applied in light transport simulation to increase MCMC efficiency:

- A simple two chain approach with a binary and uniform target function used to distribute photons in a scene [Hachisuka and Jensen 2011].
- A method using several chains, each useful at sam-pling different types of paths. However, some types of paths (e.g. reflected caustics) are not sampled well by any of the chains [Kitaoka et al. 2009].

Tempering of light transport

To maximize efficiency of parallel tempering, each auxiliary chain should have increasingly flatter target function, so it can more easily explore the path space.

Our solution:
We increase roughness of bidirectional reflectance function (BRDF) lobes at connecting vertices of the two subpaths (one from the eye/camera, one from the light source. The auxiliary chains (blue, red, green) therefore transport more energy through the connection than the main chain (black).

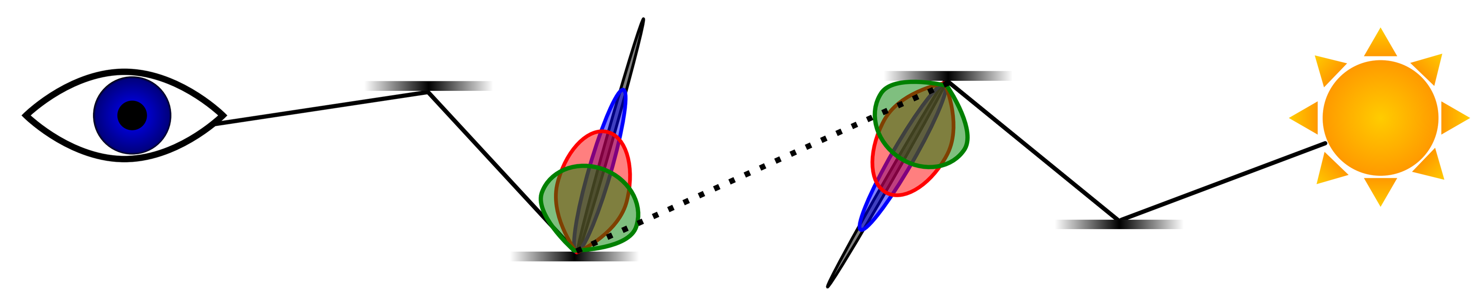

Visualization of the tempered path space for each chain:
Tempering/roughening BRDF lobes results in increasingly simpler target functions that can be easily explored and thus the auxiliary chains discover all the target function extremes.

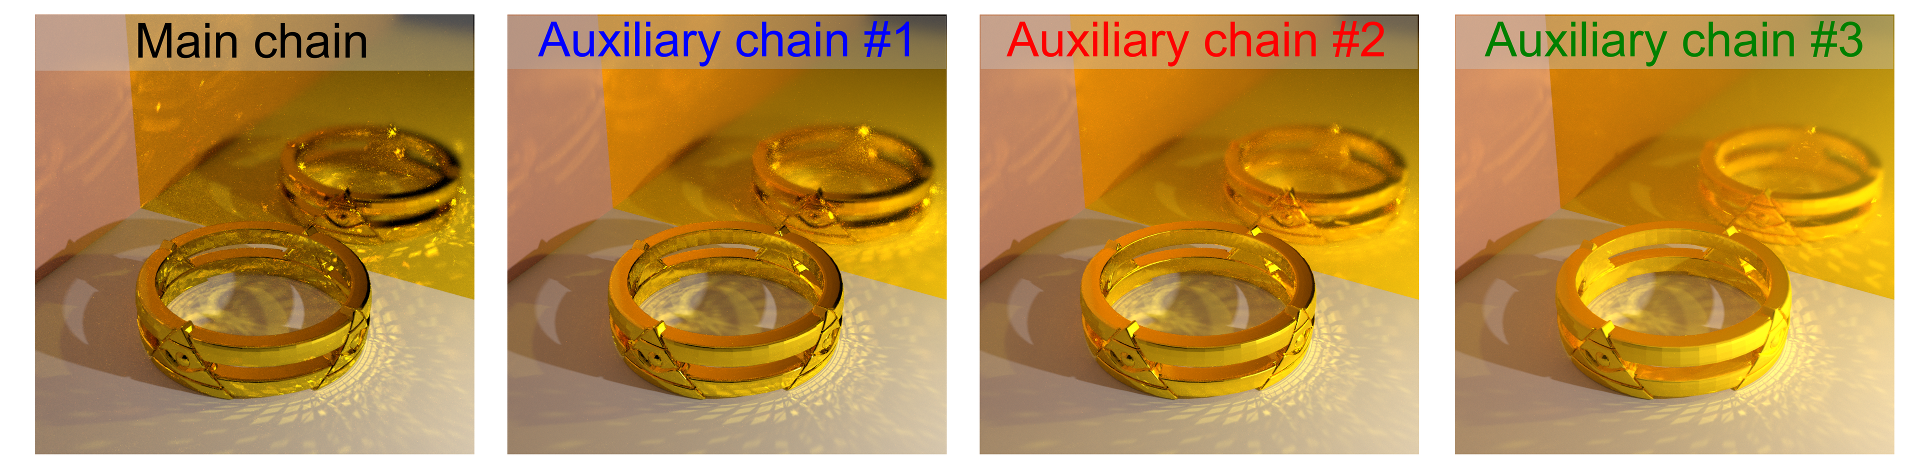

Swapping strategies

Exploration of the path space by chains with more complex target functions is increased by probabilistically swapping their samples with chains that use flatter target functions. We propose three different swapping strategies which try to maximize the probability of accepting selected swaps.

Neighbor swapping

The classic strategy - only swaps chains with the most similar target functions.

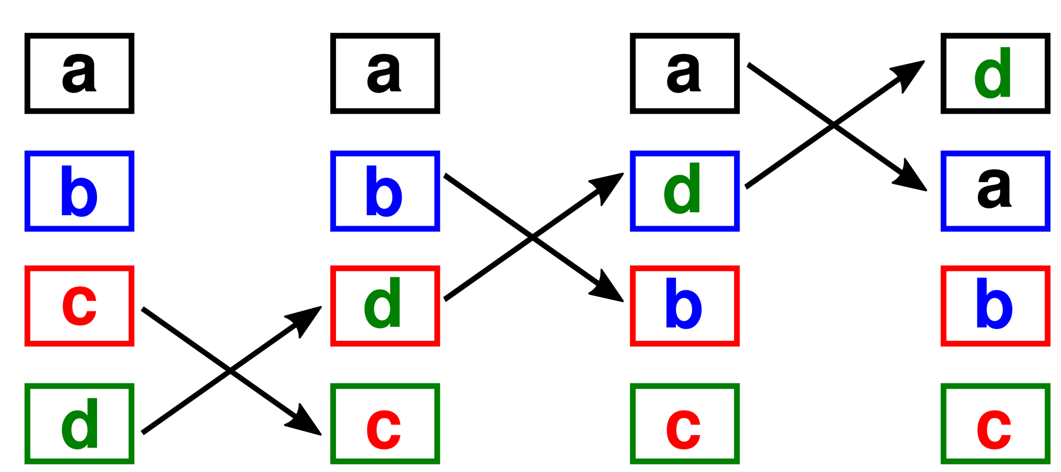

Equi-energy moves [Baragatti et al. 2012]

This strategy swaps two current samples that have the most similar target function value.

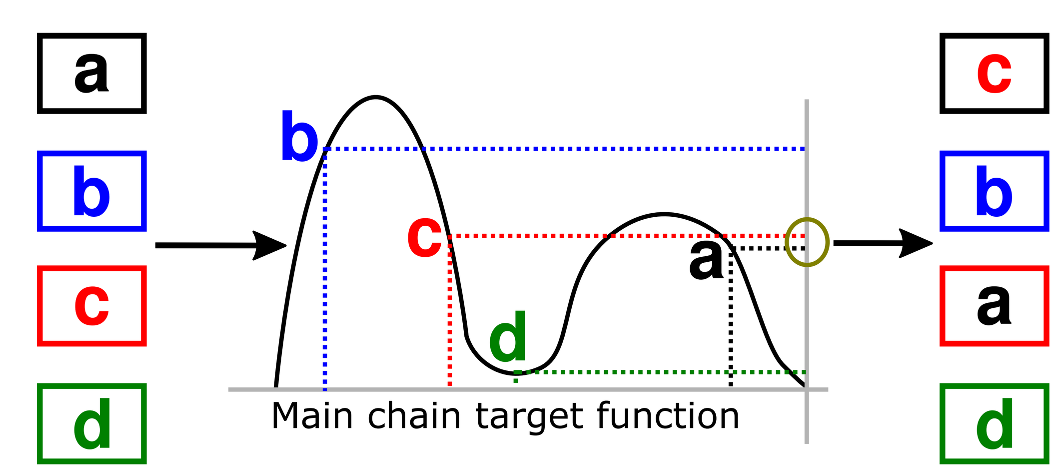

Importance-sampled permutations

Our novel strategy that importance-samples one permutation from all sample permutations.

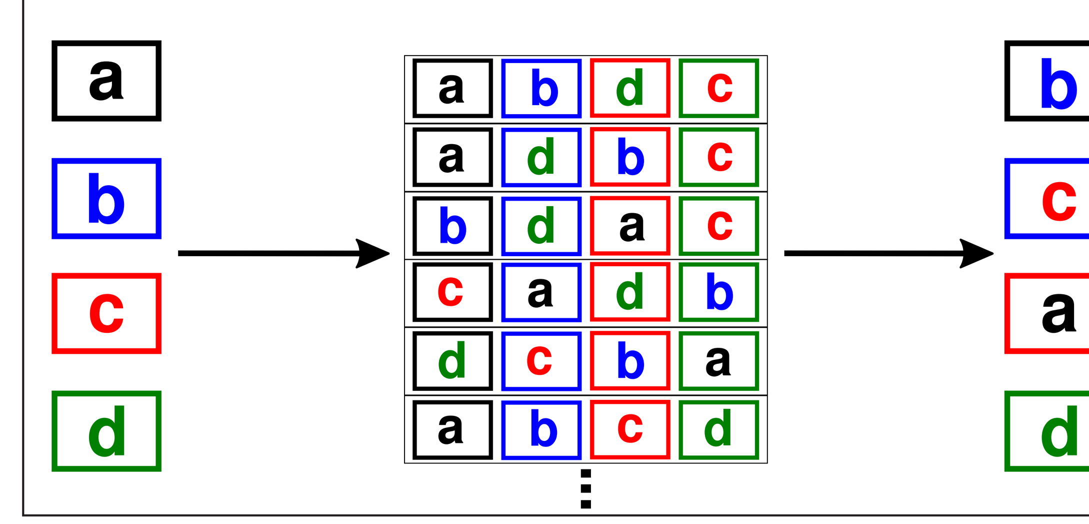

Results

Equal-time comparison of the original Multiplexed Metropolis Light Transport [Hachisuka et al 2014] (MMLT) to three variants of our algorithm (which is built on top of MMLT). Due to insufficient global exploration in original MMLT some of the transport (e.g. reflected caustics) is missing in the image, while our algorithm is able to recover it. The three variants of our algorithm perform roughly on par.

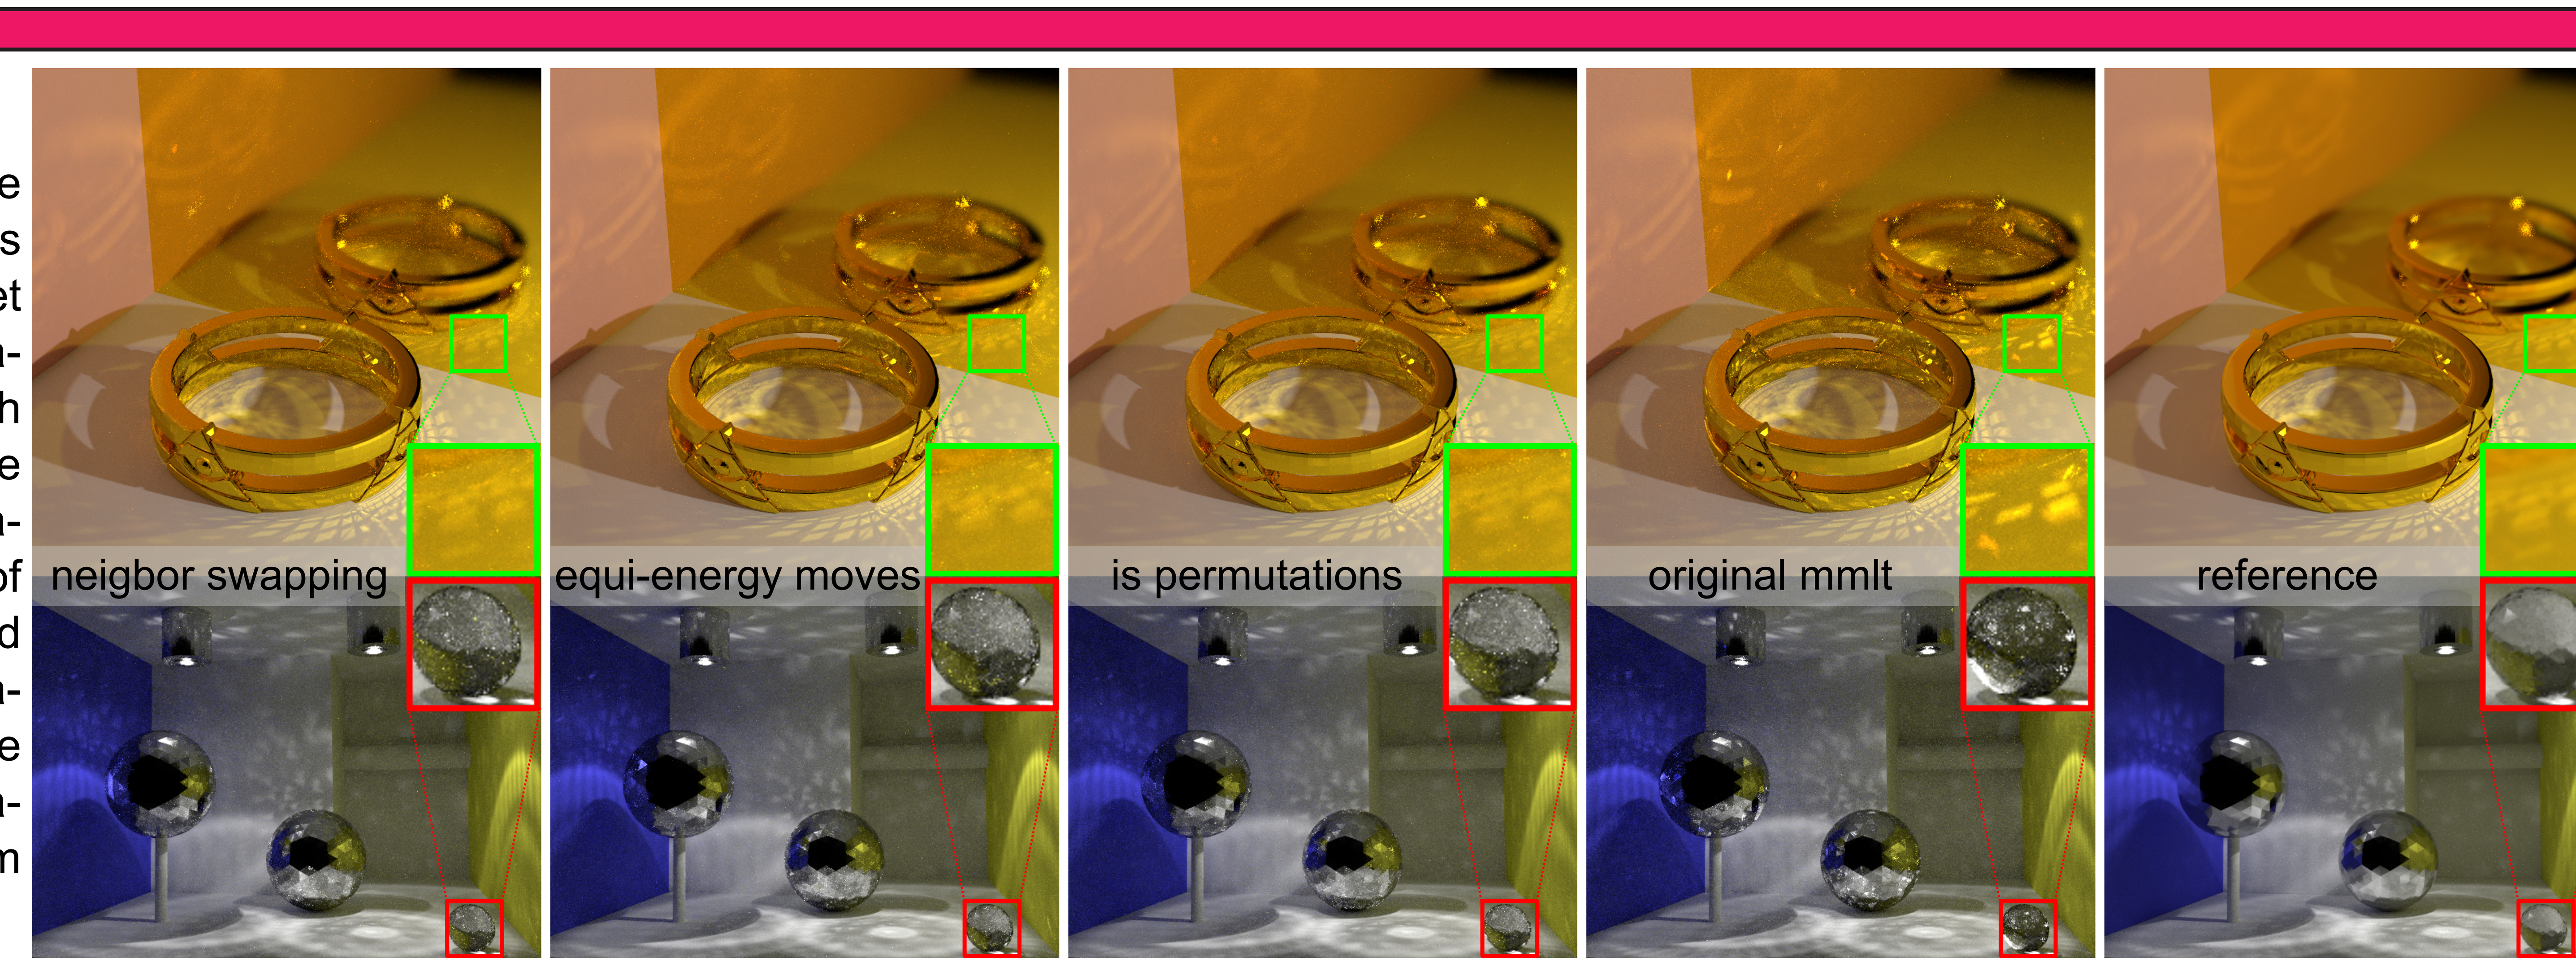

Acknowledgements: The work was supported by Charles University in Prague, project GA UK 164815, by the grant SVV–2016–260332, and by the Czech Science Foundation grant 16–18964S