Microfacet-Based Normal Mapping
for Robust Monte Carlo Path Tracing

Compatibility with Metropolis Light Transport

Vincent Schüßler[1] Eric Heitz[2] Johannes Hanika[1] Carsten Dachsbacher[1]
[1]Karlsruhe Institute of Technology [2]Unity Technologies

In this document, we explain how stochastically evaluated BSDFs could work with Metropolis light transport at the cost of some modifications made to the integrator. This discussion concerns the stochastic evaluation as computed in our plugin **normalmap_microfacet**. Note that our plugin **normalmap_microfacet_default** works straightforwardly with any light transport algorithm since it is analytic. We provide results (Fig. 1) from the Smith multiple-scattering BSDF model of Heitz et al. [2016] that is subject to the same limitations as microfacet-based normal mapping: both evaluate and sample the BRDF with random walks.

In either case, Kelemen- or Veach-style MLT, the idea is to regard the random walk inside the material as a subpath of a light path in the scene: we extend the path space to what happens inside the microsurface. With this extension, a light path can be made of "macro"-vertices collected in the scene (typically ray-triangle intersections) and of "micro-vertices" collected by the random walk in the material plugin. With this formulation of light transport the contribution and PDF of each light path can be evaluated. Indeed, we do not know the PDF of the full stochastic BSDF model but we know the PDF of each random walk: the contribution and PDF of one light path (random walk) inside the microsurface is given by the product of the micro-BSDF at each vertex and the microsurface intersection probability (which is the equivalent of the vertex area measure inside the microsurface). The Metropolis integrator can mutate light paths as usual, i.e. mutations affect macro- or micro-vertices, or both. To make this work, the main modification to be done is to make the integrator aware that vertices collected inside the material plugin should be handled in the same way as vertices collected in the scene.

Kelemen Metropolis In a Kelemen-style or primary sample space Metropolis implementation, the transition probabilities are computed as the ratio of path weights (which are measurement contribution divided by vertex area measure PDF). These weights are computed by the path sampler, as the path is constructed, and usually never involve computation of a PDF explicitly. Running a fresh microsurface random walk (to stochastically evaluate the BSDF or to facilitate sampling) can be seen as sampling in an extended primary sample space which holds the random numbers necessary for these extra random walks. Since we do not explicitly store and mutate these extra random numbers, we implicitly perform a large step perturbation on them in every step of the Markov chain, i.e. the random numbers are chosen uniformly random in the full interval $[0, 1)$ every time. The other extreme variant, always leaving the random walk fixed for a small step mutation, is also possible. With this algorithm, primary sample space variants of MLT work transparently with a stochastic evaluation of the BSDF.

Veach Metropolis A Markov chain running mutations in path space, i.e. evaluating the Metropolis-Hastings transition probabilities in vertex area measure explicitly, may have to compute the PDF of sampling a certain direction and would divide it out explicitly from the measurement contribution function. Fortunately, the evaluation and PDF of a single random walk can be evaluated in closed form, and so can be the evaluation and contribution of any light path in the scene. In practice, we make a vertex always use the same microsurface random walk. The microsurface random walk can be derived with a pseudo random number sequence which is seeded with a deterministic hash that depends on the path id and the vertex id. The other random walks can be explored by the next mutation step, i.e. the choice of the random walk at each vertex becomes one mutation strategy among others.

References

Eric Heitz, Johannes Hanika, Eugene d'Eon, and Carsten Dachsbacher. 2016. Multiple-scattering Microfacet BSDFs with the Smith Model. *ACM Trans. Graph. (Proc. SIGGRAPH)* 35, 4 (July 2016), 58:1–58:14. (document), 1

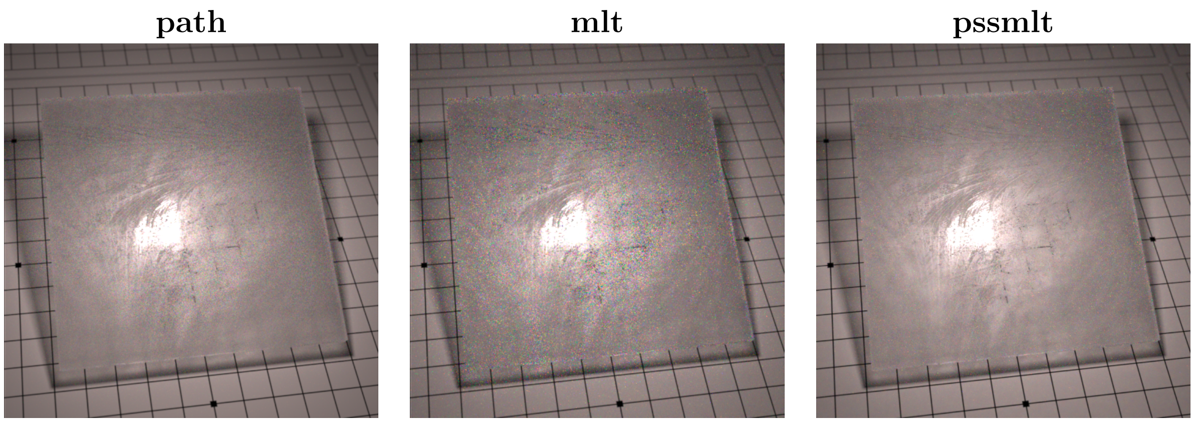

Figure 1: **Results using the multiple-scattering BSDF model of Heitz et al. [2016].** A multiple-scattering Smith microsurface with dielectric material, rendered with (from left to right) path tracing with next event estimation, path space Metropolis light transport (multichain perturbation), and primary sample space Metropolis light transport, using path tracing and next event estimation as underlying path construction scheme. All images use 8k samples/pixel. The images demonstrate that it is possible to combine stochastically evaluated BSDFs with complex path sampling strategies.