Bidirectional Light Transport with Vertex Merging

Iliyan Georgiev*
Saarland University
Intel VCI Saarbrücken

Jaroslav Křivánek†
Charles University, Prague

Philipp Slusallek‡
DFKI, Saarland University
Intel VCI Saarbrücken

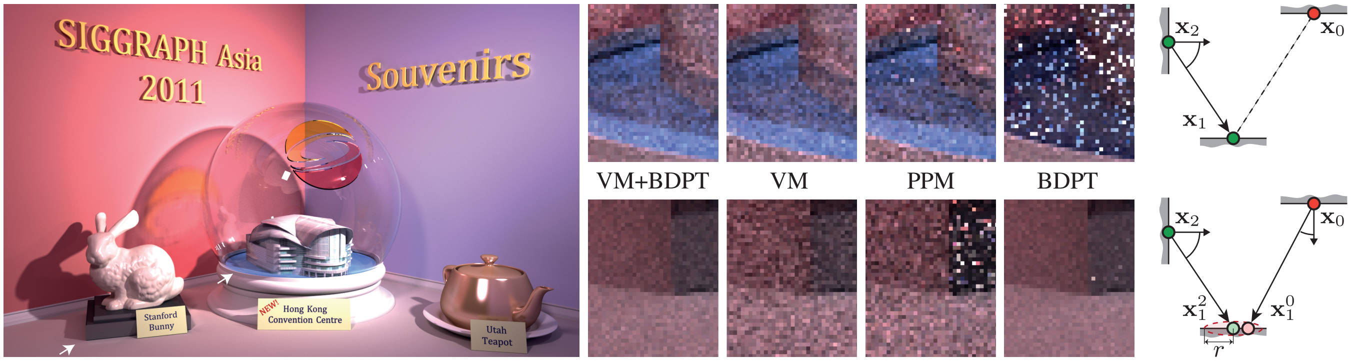

Figure 1: *Left: The SIGGRAPH Souvenirs scene reference image. Middle: Insets showing the quality achieved by 4 different algorithms in the same time (200 seconds): Combined vertex merging with bidirectional path tracing (VM+BDPT), vertex merging only (VM), progressive photon mapping (PPM), bidirectional path tracing (BDPT). The new VM+BDPT algorithm combines various path sampling techniques to produce the image with the lowest overall error. Right: A path segment generated by a vertex connection (top) and by vertex merging (bottom).*

Abstract

We present vertex merging – a bidirectional path sampling technique for Monte Carlo light transport integration. Vertex merging is simple and more computationally efficient for specular-diffuse-specular effects than the currently available techniques in bidirectional path tracing. It brings the advantages of photon mapping to the path integral framework, while avoiding the concept of density estimation altogether. This makes it possible for the first time to quantitatively reason about the efficiency of two rendering approaches that have been historically considered conceptually different. The practical result is a combined bidirectional rendering algorithm that efficiently handles a wide variety of lighting conditions, ranging from direct illumination and diffuse inter-reflections to the notoriously problematic reflected caustics. This algorithm also has a higher order of convergence than progressive photon mapping.

1 Light Path Sampling

Veach [Veach 1997] formulated the rendering problem mathematically as a radiance measurement function integrated over the space of all light transport paths. This tidy *path integral framework* made it possible to combine different Monte Carlo light transport estimators in an efficient way using multiple importance sampling (MIS). Veach then developed bidirectional path tracing (BDPT) as a family of path sampling techniques whose corresponding estimators are weighted with the power heuristic [Veach 1997]. This heuristic assumes that higher probability density function (pdf) values result in lower variance, and weights estimators proportionally to the pdf of their corresponding sampling technique.

The two basic building blocks of BDPT are the unbiased *unidirectional sampling* and *vertex connection* techniques. Unidirectional sampling constructs a path connecting a light source with the camera by performing a random walk from either end until termination. Vertex connection joins the endpoints of a camera and a light subpaths deterministically by a ray, saving one random sampling step.

*e-mail: georgiev@cs.uni-saarland.de
†e-mail: jaroslav.krivanek@mff.cuni.cz
‡e-mail: slusallek@dfki.de

Vertex connection works best if the two vertices are far apart and both have low frequency BSDFs. Unidirectional sampling is in contrast superior for specular paths, and inferior for diffuse inter-reflections. Therefore, specular-diffuse-specular (SDS) vertex sequences are ill-suited for both techniques, and result in low sampling pdfs. Such paths occur, e.g., when looking at a car's interior or at a glass-enclosed object from the outside, as shown in Figure 1.

Photon mapping has been shown to efficiently handle SDS interactions in a consistent way [Hachisuka et al. 2008]. Unfortunately, it is inefficient for diffuse illumination and has a lower order of convergence than the unbiased BDPT estimators [Knaus and Zwicker 2011]. Furthermore, the flux density estimator has been derived in a framework that cannot be easily mapped to the path integral. These different theoretical foundations have prevented the rigorous comparison, and hence an efficient combination, of the two algorithms.

2 Vertex Merging

Our new path sampling technique is motivated by the observation that diffuse BSDF sampling can often bring additional variance, whereas sampling specular interactions is usually a good importance sampling strategy. We therefore want to be able to construct SDS paths by performing directional sampling only at specular surfaces. This necessitates a bidirectional approach.

Vertex merging builds approximate paths by virtually welding the endpoints of a camera and a light sub-paths into one, if they are within distance r. This is illustrated in the bottom right of Figure 1.

Assume we have already created $\mathbf{x}_0$ and $\mathbf{x}_2$ with cumulative subpath pdfs p^0 and p^2, as well as $\mathbf{x}_1^2$ from $\mathbf{x}_2$. We now sample a ray from $\mathbf{x}_0$. If its hit point $\mathbf{x}_1^0$ lies in a circular region around $\mathbf{x}_1^2$ with radius r, then we generate the final path by merging $\mathbf{x}_1^0$ and $\mathbf{x}_1^2$ into $\mathbf{x}_1 \equiv \mathbf{x}_1^2$. The full path pdf is then $p = p^0 p(\mathbf{x}_1) p^2$, where

$$p(\mathbf{x}_1) = p(\mathbf{x}_1^0, \mathbf{x}_1^2) = p(\mathbf{x}_1^2) p(\mathbf{x}_1^0 | \mathbf{x}_1^2)$$
$$\approx\; p_\sigma(\mathbf{x}_2 \to \mathbf{x}_1^2) \frac{\cos\theta^{1^2 \to 2}}{||\mathbf{x}_2 - \mathbf{x}_1^2||^2}\, p_\sigma(\mathbf{x}_0 \to \mathbf{x}_1^0) \frac{\cos\theta^{1^0 \to 0}}{||\mathbf{x}_0 - \mathbf{x}_1^0||^2} \pi r^2,$$

where πr^2 is the area of the region, p_σ denotes a solid angle pdf, and $\theta^{i \to j}$ the angle between the normal at $\mathbf{x}_i$ and vector $\mathbf{x}_i \to \mathbf{x}_j$.

To apply vertex merging in practice, we trace one path from the camera and one from a light source, and then merge their endpoints. Similarly to the BDPT techniques, we can reuse the sub-paths by merging any camera vertex with any light vertex.

Having a path constructed with our new technique and its pdf, we can obtain a pixel estimator by evaluating the path's radiance contribution. For the example in Figure 1 the estimator expands to

$$\widetilde{L} = \frac{1}{\pi r^2} \underbrace{\frac{L^{0\to1}\cos\theta^{0\to1}}{p^0 p_\sigma(\mathbf{x}_0 \to \mathbf{x}_1^0)}}_{\Delta\Phi_1^0} f_r(\mathbf{x}_0,\mathbf{x}_1,\mathbf{x}_2) \underbrace{\frac{W^{2\to1}\cos\theta^{2\to1}}{p^2 p_\sigma(\mathbf{x}_2 \to \mathbf{x}_1^2)}}_{\Delta W_1^2}.$$

Surprisingly, $\widetilde{L}$ is equivalent to the photon mapping estimator at $\mathbf{x}_1^2$, with $\Delta\Phi_1^0$ being the incident photon flux at $\mathbf{x}_1^0$ and ΔW_1^2 the cumulative importance, or "path throughput", up to vertex $\mathbf{x}_1^2$.

2.1 Comparison to BDPT

Avoiding the explicit use of the concept of density estimation, we have reformulated photon mapping as a path sampling technique, where we allow identification of opposing path vertices within a small neighborhood. The path integral framework now allows us to rigorously compare its efficiency against the BDPT techniques.

Consider the two extreme cases in Figure 2. For diffuse interactions, vertex connection (VC) can have orders of magnitude higher pdf than both unidirectional sampling (US) and vertex merging (VM). The reason is that VC saves the random directional sampling, which has a very low pdf due to the BSDF sampling and geometric relation between $\mathbf{x}_0$ and $\mathbf{x}_1$. For SDS paths, VC has zero pdf, while US and VM can in fact be shown to have the same pdf for the case in Figure 2 right, if the light source area equals πr^2.

We observe that VM, and thus photon mapping, is not an intrinsically more robust sampling technique than US and VC. However, its strength is computational efficiency. Since path joining is as cheap as neighborhood checking, VM allows the reuse of any previously generated light (sub-)paths in the vicinity at the cost of a single range search. Therefore, in cases where other techniques have low pdfs, VM can result in a much lower error estimate due to its efficient brute-force reduction of variance, which is inversely proportional to the *total* number of light paths started from light sources [Knaus and Zwicker 2011]. The most prominent example for such cases are the SDS paths.

3 A Combined Algorithm

Thanks to the new formulation as a path sampling technique, we can plug vertex merging directly into the multiple importance sampling (MIS) pixel estimator together with the BDPT techniques. With the observations above, we can expect to benefit from an efficient combination that is made possible by the power heuristic.

Since path sampling is expensive, we want to amortize as much of this effort as possible among the techniques. Fortunately, vertex merging allows for an implementation that comes at little extra cost.

Rendering is done in two stages. We first trace a set of paths from the light sources and store their vertices by additionally building a range search data structure over them, e.g. a kd-tree or a grid. We will then use this data for vertex connections and merging.

In the second stage, we perform slightly extended path tracing. As we do the random walks from the camera, we perform connections between the current vertex and the light sources and a number of light vertices. Additionally, we perform a range search at each camera vertex to merge all light vertices in the local r-neighborhood.

We make an additional optimization in the computation of the MIS weights. During each random walk (from either end), we keep track of three floating point numbers that store cumulative forward and reverse sampling pdfs along the path. Whenever a technique builds an estimator, e.g. when we hit a light/camera or connect or merge two vertices, the cumulative data stored at the endpoints of the two sub-paths being joined is sufficient to compute the MIS weight of the used technique w.r.t. all other possible ways of generating the full path. This brings performance speed-up by removing the need for traversing back the vertices of the whole path only to compute a weight. This also removes the need for storing the camera vertices.

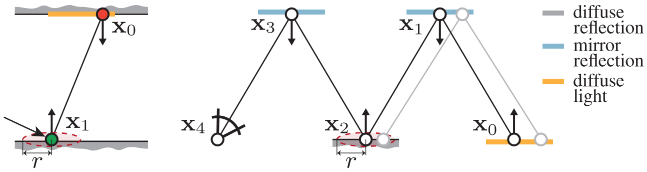

Figure 2: *Left: For diffuse (sub-) paths, vertex connection (VC) can have orders of magnitude higher pdf than unidirectional sampling (US) and vertex merging (VM). Right: For SDS paths, VC has zero pdf, and US and VM can have pdfs of the same order, but VM can very efficiently reuse the light subpaths created for other pixels.*

4 Results and Discussion

Figure 1 shows a rendering setup that contains various illumination effects, and compares the quality achieved by four methods in the same rendering time. The scene is problematic for bidirectional path tracing (BDPT) particularly due to the exclusive presence of SDS paths on the glass ball. The dominating diffuse illumination, on the other hand, makes it difficult for progressive photon mapping (PPM). Vertex merging (VM) alone outperforms PPM in terms of quality, since it operates on all possible points along the path, combining all estimates with multiple importance sampling. Nevertheless, it still cannot handle diffuse illumination as well as BDPT. The new combined VM and BDPT algorithm (VM+BDPT) takes the best of both worlds, with the power heuristic automatically finding a good mixture of techniques for each individual light path. The BDPT image is slightly less noisy than VM+BDPT in the lower inset, as BDPT alone can take a bit more samples within the given time. The four comparison images took 200 seconds to render on an Intel Core i7 860 CPU at 750×600 resolution. (The full images are provided as supplemental material.) We let the reference image render for 5 hours at 1500×1200 resolution.

A significant advantage of the VM+BDPT algorithm over PPM is that its variance vanishes to the order of $O(1/N)$ in contrast to the rate of $O(1/N^\alpha)$, $\alpha \in (0;1)$, for PPM [Knaus and Zwicker 2011]. For an intuitive explanation of this behavior, note first that the pdf of the VM technique is proportional to r^2. As r approaches zero (to ensure consistency), so does its weight in the power heuristic. Conversely, the relative weights of all unbiased BDPT techniques converge to one, and their order of convergence is $O(1/N)$. As a conclusion, the VM technique can bring tremendous variance reduction initially, but its efficiency diminishes over time.

References

HACHISUKA, T., OGAKI, S., AND JENSEN, H. W. 2008. Progressive Photon Mapping. *ACM Trans. Graph. (SIGGRAPH Asia Proceedings) 27*.

KNAUS, C., AND ZWICKER, M. 2011. Progressive Photon Mapping: A Probabilistic Approach. *ACM Trans. Graph. 30*.

VEACH, E. 1997. *Robust Monte Carlo Methods for Light Transport Simulation*. PhD thesis, Standford Univeristy.