The Visual Computer manuscript No.
(will be inserted by the editor)

Abhijeet Ghosh · **Wolfgang Heidrich**

Correlated Visibility Sampling for Direct Illumination

Abstract State-of-the-art importance sampling strategies for direct illumination take into account the importance of the incident illumination, as well as the surface BRDF. Hence, these techniques achieve low variance in unoccluded regions. However, the resulting images still have noise in partially occluded regions as these techniques do not take visibility into account during the sampling process.

We introduce the notion of *correlated visibility sampling*, which considers visibility in partially occluded regions during the sampling process, thereby improving the quality of the shadowed regions. We aim to draw samples in the partially occluded regions according to the triple product of the incident illumination, BRDF and visibility using Monte Carlo sampling followed by Metropolis sampling.

Keywords Monte Carlo Techniques · Ray Tracing

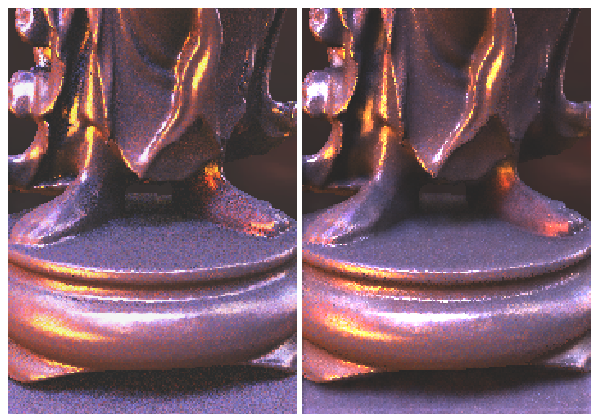

Fig. 1 Buddha model (Phong BRDF $s = 50$, $k_s = k_d = 0.5$) in Grace Cathedral EM. Left: Bidirectional importance sampling, 20 samples/pixel. Right: Correlated visibility sampling, 16 bidirectional samples (1^{st} pass) and 16 Metropolis samples per unoccluded sample (2^{nd} pass). Rendering times are identical (24 seconds).

1 Introduction

Image-based representations for illumination, such as environment maps, textured area lights, and light fields, have become very popular in recent years as these images can capture complex real-world illumination that is difficult to represent in other forms.

The use of a good sampling strategy for illumination is paramount when integrating image-based lighting, such as environment maps, into a rendering system. This is because direct illumination in the form of high dynamic range (HDR) environment maps can have high frequency detail. The problem of efficient sampling of the illumination is compounded when the scene contains materials with high frequency BRDFs. Several researchers have recently worked on this problem, by either combining samples drawn independently according to the importance of the illumination and the BRDF [31], or more recently, by drawing samples from the *product* distribution of the illumination and the BRDF [2] [3] [30]. These approaches produce high quality images with a small number of samples in unoccluded regions. However, the resulting images still have noise in partially occluded regions as these techniques do not take visibility into account in the sampling process (Figure 1, left).

This paper introduces *correlated visibility sampling*, a method that additionally takes visibility into account in the sampling process for partially occluded regions, thereby improving the quality of the shadowed regions (Figure 1, right). The aim of this technique is to develop an efficient means of drawing samples from the triple product of the incident illumination, BRDF and visibility, which we achieve by employing the Metropolis-Hastings (MH) algorithm [18]. We describe two variants of the method, one which is unbiased, and a more efficient one that is consistent, but may exhibit a small amount of bias.

Our solution is a two-step approach. In the first step, energy estimates for each pixel are created using samples drawn from the bidirectional importance (product distribu-

A. Ghosh, W. Heidrich
The University of British Columbia
2366 Main Mall, Department of Computer Science
Vancouver, B.C., V6T 1Z4, Canada.
Tel.: +1-604-822-4326
Fax: +1-604-822-8989
E-mail: {ghosh,heidrich}@cs.ubc.ca

tion) of the incident illumination and the surface BRDF (Section 3). This estimate is built using the sampling-importance resampling (SIR) algorithm, as discussed by Burke et al. [2]. We create a visibility mask and mark pixels for which one or more of the visibility tests failed, i.e., pixels which are partially occluded. In the second step, Metropolis sampling is started for the partially occluded pixels in order to locally explore the shadowed regions more extensively (Section 4). If desired, any image-space operation such as dilation can be applied to the visibility mask. Our approach has the following benefits:

– The energy estimates from the first round of sampling are created with a small number of visibility tests. For unoccluded regions, this small number of samples is sufficient for providing a good estimate of the integral.
– By employing the visibility mask, visibility tests can then be restricted in the second phase to the partially occluded regions where more samples are required in order to achieve low variance.
– The sampling in the second phase can exploit correlation in the energy estimates of neighboring pixels as a powerful tool for variance reduction.
– Metropolis sampling in the second phase is started only from those bidirectional samples that passed visibility tests in the first phase as these are the valid samples according to the target distribution. Markov chains are started from an *unbiased* Monte Carlo estimate and hence, have no startup bias.

The rest of this paper is structured in the following manner. Section 2 reviews some of the relevant work in Monte Carlo sampling from environment maps as well as product distributions and Metropolis sampling for global illumination. Section 3 gives an overview of the bidirectional sampling approach which is employed in the first phase of our algorithm. Section 4 presents the idea of correlated visibility sampling which we employ in the second phase of our solution. We conclude with results and a discussion in Sections 5 and 6.

2 Related Work

The computation of the direct illumination in a scene is a costly task in all rendering systems, both global and local. This task is complicated in presence of complex real-world light sources such as environment maps and other image-based representations. Much effort has focused on the development of efficient techniques for completing this task.

2.1 Sampling from Environment Maps and BRDFs

Illumination from environment maps has been a topic of much recent research. Most of this work focuses on interactive applications and therefore uses expensive precomputation [8–11]. In some recent work, the illumination and/or the BRDF are projected into finite bases such as spherical harmonics (e.g., [21, 22, 25]) and wavelets [19].

Other researchers have used importance sampling techniques to distribute samples according to the energy distribution in the environment map, either by using point relaxation schemes [5, 13] based on Lloyd's clustering algorithm [16] or by using an efficient hierarchical Penrose tiling scheme [20].

Agarwal et al. [1] introduced a sampling method for environment maps taking into account both the energy distribution in the environment map and the solid angle separating the samples. In this way, close clustering of environment map samples is avoided, which reduces redundant shadow tests.

In the context of stippling, Secord et al. [23] described an algorithm for computing and inverting the cumulative density function (CDF) based on image intensities. Inversion of CDFs is also used by Lawrence et al. [15] to sample from environment maps. This is a simple and efficient method, a variant of which we use in our work for drawing samples from environment maps.

Importance sampling from the BRDF is a common operation, though the exact mechanics depend on the specific representation used. Simple analytical models such as diffuse, Phong, or generalized cosine models can be sampled analytically (see e.g., [24]). For tabulated BRDFs, kd-tree representations [17] and more recently, factored representations [14] have been used for efficient importance sampling. For procedural shaders, cosine lobe approximations have been used for importance sampling [26]. In our work, we use only Phong and diffuse reflection models. However, our method could easily be extended to incorporate more sophisticated materials using any of the above methods.

2.2 Multiple Sampling Approaches

Veach & Guibas [31] proposed multiple importance sampling as an effective variance reduction technique. They combined different sampling distributions such as illumination and BRDF distributions in an optimal manner using their proposed *balance* heuristics. However, methods that directly sample from the product distribution generally reduce variance further than is done by multiple importance sampling [2, 3].

Szecsi et al. [29] sample the unoccluded illumination using correlated sampling and the difference due to visibility using importance sampling. This method generally performs well in fully visible regions, but rather poorly in occluded or partially occluded regions, since the sampling of visibility does not follow a special sampling pattern.

2.3 Sampling from Product Distributions

Burke et al. [2] introduced the notion of *bidirectional sampling* which takes into account the energy of incident illumination as well as the BRDF in the sampling process. They present two Monte Carlo algorithms for sampling from the product distribution - one based on rejection sampling and the other based on sampling-importance resampling (SIR), the latter of which is also used by Talbot et al. [30]. In our work, we use their SIR algorithm during our first phase of bidirectional sampling.

Clarberg et al. [3] present a technique for efficiently sampling the product of the illumination and the BRDF using a hierarchical wavelet representation. Their method is very efficient for tabulated BRDFs but requires significant precomputation for environment maps. Lawrence et al. [15] present an approach for compressing cumulative distribution functions for efficient inversion and they apply it to sampling from many precomputed environment map PDFs for different surface orientations, which is a step towards approximating the product distribution.

2.4 Metropolis Sampling for Global Illumination

Veach & Guibas [32] first applied Metropolis sampling to the problem of image synthesis and developed a general, robust and unbiased algorithm called Metropolis Light Transport (MLT) that was well suited for hard cases for sampling because of its localized exploration and path re-usage properties. Fan et al. [7] recently applied the Metropolis algorithm for efficiently sampling coherent light paths for photon mapping.

Cline et al. [4] presented an efficient unbiased method to solve correlated integral problems with a hybrid algorithm that uses Metropolis sampling-like mutation strategies in a standard Monte Carlo integration setting, overcoming the startup bias problem of MLT. They apply energy redistribution over the image plane to reduce variance of path tracing for global illumination. Our work is similar in spirit to theirs in the sense of using initial Monte Carlo sampling followed by Metropolis sampling, except that we apply this to direct illumination with a specific focus on efficient exploration of visibility in partially occluded regions.

3 Bidirectional Importance Sampling

As mentioned in the introduction, we create energy estimates for each pixel in our first round of sampling using bidirectional importance sampling [2]. This sampling approach takes both the energy distribution in the illumination and the reflectance of the BRDF into account. The rationale for using this technique for first round sampling is that it requires very few visibility tests for achieving good quality in unoccluded regions.

Consider the direct illumination at a point for a given observer direction ω_r:

$$L_r(\omega_r) = \int_\Omega f_r(\omega_i \to \omega_r)\cos\theta_i L_i(\omega_i)V(\omega_i)d\omega_i, \tag{1}$$

with L_i denoting the incident illumination from an environment map, f_r representing the BRDF, and V being the binary visibility term.

The aim of bidirectional sampling is to perform importance sampling using the *product* of the incident light distribution and the BRDF as the importance function:

$$p(\omega_i) := \frac{f_r(\omega_i \to \omega_r)\cos\theta_i L_i(\omega_i)}{\int_\Omega f_r(\omega_i \to \omega_r)\cos\theta_i L_i(\omega_i)d\omega_i}. \tag{2}$$

Observe that the normalization term in the denominator is the direct illumination integral with the visibility term $V(\omega_i)$ omitted. In other words, this term is the exitant radiance in the absence of shadows. Burke et al. refer to it as L_{ns} "radiance, no shadows" [2]:

$$L_{ns} := \int_\Omega f_r(\omega_i \to \omega_r)\cos\theta_i L_i(\omega_i)d\omega_i. \tag{3}$$

If sample directions $\omega_i^{(j)} \sim p(\omega_i),\ j = 1,...,N$, are drawn according to the product distribution in Equation 2, then Equation 1 can be estimated as $L_{N,p}$, where

$$\begin{aligned} L_{N,p}(\omega_r) &= \frac{1}{N}\sum_{j=1}^{N}\frac{f_r(\omega_i^{(j)} \to \omega_r)\cos\theta_i^{(j)}L_i(\omega_i^{(j)})V(\omega_i^{(j)})}{p(\omega_i^{(j)})}; \\ &= \frac{L_{ns}}{N}\sum_{j=1}^{N}V(\omega_i^{(j)}). \end{aligned} \tag{4}$$

$L_{N,p}$ is referred to as the *bidirectional estimator* for the direct illumination integral.

Bidirectional sampling is implemented as a two-step approach: initially samples are created according to either the BRDF alone or the environment map alone, and then these samples are adjusted to be proportional to the product distribution. The adjusted samples are then used for visibility testing. Two Monte Carlo algorithms are presented in [2] to achieve bidirectional sampling – one based on rejection sampling and the other based on sampling-importance resampling (SIR) [27]. We employ the SIR algorithm in this work for our first round of sampling due to the deterministic nature of its execution time compared to rejection sampling.

Note that the variance of the bidirectional estimator for the reflected radiance is proportional to the variance in the visibility function. This is an improvement over sampling techniques that only consider either the illumination or the BRDF in the sampling process. This is because the variance of these techniques depends in addition on the variance in the function that they do not sample from, BRDF or illumination respectively. However in regions with complex visibility, estimates with bidirectional sampling will still suffer from considerable variance. That is the problem we attempt to solve in this work using our correlated sampling algorithm.

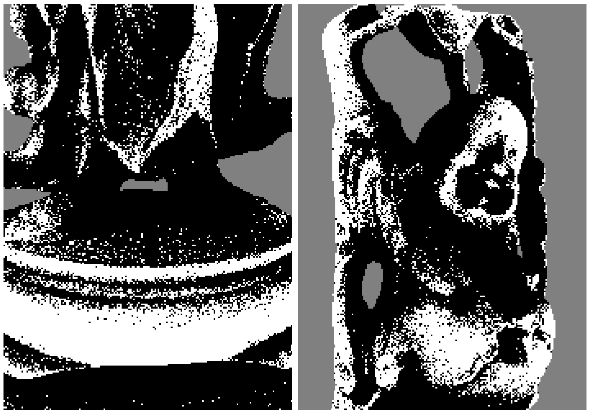

Fig. 2 Visibility masks for the images presented in Figures 1 and 5. The white pixels correspond to unoccluded pixels at the end of first round of bidirectional sampling while the black pixels correspond to the partially occluded pixels.

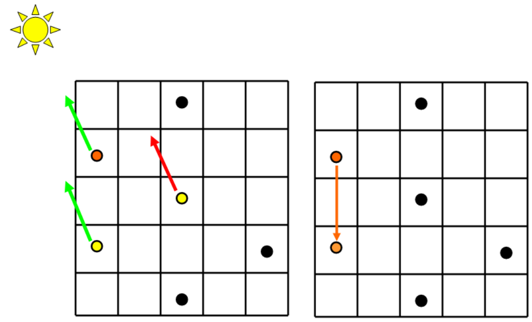

Fig. 3 Lens perturbation within a 5×5 transition tile. Left: The source (orange) pixel selects two other (yellow) pixels within the transition tile for energy transfer. Right: Only one pixel is selected for energy transfer based on visibility test in the same direction. Green arrows refer to unoccluded light directions, red arrows to occluded ones.

4 Correlated Visibility Sampling

In this work, we propose a correlated sampling approach that takes visibility into account in addition to the incident illumination and the surface reflectance. This is a two-step process: we initially create energy estimates for each pixel using bidirectional importance sampling as discussed in Section 3. We create a visibility mask and mark pixels for which one or more of the visibility tests failed, i.e., pixels that are partially occluded (Figure 2). In the second step, we employ the Metropolis-Hastings algorithm in order to locally explore visibility in the shadowed regions more extensively. Metropolis sampling is only started for the partially occluded pixels, thereby avoiding unnecessary visibility tests in unoccluded regions.

Given a non-negative function f, the MH algorithm generates a series of samples $X = \{x_1, x_2, ..., x_n\}$ from a distribution proportional to f, which is also referred to as the target distribution, without requiring to normalize f and invert the resulting PDF. It is thus applicable to a wide variety of sampling problems and was first applied in computer graphics by Veach & Guibas [32] to the problem of image synthesis. Given a current sample x, the next sample x' in the sequence is generated by randomly mutating x and then accepting or rejecting the mutation. The mutations are accepted or rejected in such a way that the samples converge to the target distribution. For a description of the MH algorithm, we refer the reader to the chapter on Metropolis Sampling by Pharr in the SIGGRAPH 2004 course notes [6].

The MH algorithm generally suffers from a *startup bias* as the initial samples in the sequence are not drawn according to the target distribution and thus need to be discarded. Despite the startup bias, integral estimates according to the MH algorithm are asymptotically unbiased as long as *detailed balance* is maintained [32]. Detailed balance defines an *acceptance probability* of a mutation strategy:

$$a(x \to x') = \min\{1, \frac{f(x') \cdot T(x' \to x)}{f(x) \cdot T(x \to x')}\}, \tag{5}$$

where x is the current sample and x' is the mutated sample, $f(x)$ is the function being integrated and $T(x \to x')$ is the cumulative *transition probability* of mutating from x to x'. Note that the acceptance probability accounts for changes in surface orientation and surface BRDF from one pixel to another, for example between the diffuse ground plane and the specular Buddha model in Figure 1.

Since we begin our Metropolis sampling from an unbiased Monte Carlo estimate, our method does not suffer from startup bias. However, in general, a small amount of bias may originate from using samples in unoccluded regions for both estimating the direct illumination, and for making a decision about entering the correlated sampling stage. This issue is discussed in more detail in Section 4.1

We employ *lens perturbation* as the mutation strategy for our algorithm. Since there is correlation in the visibility of points in neighboring pixels, using this strategy to transfer energy of samples $\omega_{i,x}$ to neighboring pixels x' can be an effective means of reducing variance. We partition the image plane into 5×5 tiles (Figure 3) for lens perturbation and carry out mutations only between the partially occluded pixels within each tile. First a mutation of a valid unoccluded sample (obtained from first round of bidirectional sampling) is proposed. Visibility is then sampled in the same direction (for environment map illumination) for the pixel that the sample gets mutated to. If the visibility test passes, the mutation is accepted with a probability a, else it is rejected. If the mutation is accepted, energy is transfered from pixel coordinate x to x'.

In our case, the cumulative transition probability $T(x \to x')$ needs to account for both the probability of choosing a neighboring pixel x' for transition from pixel x, which we will call $t(x \to x')$, and for the probability of choosing the same sample direction ω_i to sample illumination at the two pixels according to bidirectional importance p:

$$T(x \to x') = t(x \to x') \cdot p(\omega_{i,x'})$$

By restricting mutations to happen only within 5×5 tiles, we ensure that every partially occluded pixel has the same number of neighbors for energy transfer. This ensures that $t(x \to x') = t(x' \to x)$.

However, imposing fixed transition tiles could potentially lead to block artifacts at the tile boundaries. Hence, in practice we employ a *moving* tile mechanism for transition centered around the current pixel x. For example, instead of performing $C = 16$ path mutations on a single tile, we chose 16 different partitions of the image plane into 5×5 tiles with different offsets, and perform one mutation each. Following the argument from above, this yields an estimate for every pixel and each of the C tile offsets. The total estimate for one pixel is then computed as an average of $N \cdot C$ transitions from the individual tile offsets, which does not introduce bias.

The acceptance probability of the above mutation strategy then reduces to:

$$a(x \to x') = \min\{1, \frac{f(x') \cdot p(\omega_{i,x})}{f(x) \cdot p(\omega_{i,x'})}\}, \tag{6}$$

where

$$f(x) = f_r(\omega_{i,x} \to \omega_{r,x}) \cos\theta_{i,x} L_i(\omega_{i,x}),$$

since the visibility term $V(\omega_{i,x}) = 1$, and

$$p(\omega_{i,x}) = f_r(\omega_{i,x} \to \omega_{r,x}) \cos\theta_{i,x} L_i(\omega_{i,x}) / L_{ns,x},$$

where $L_{ns,x}$ is the unoccluded radiance in the viewing direction given in Equation 3. The numerator of $p(\omega_{i,x})$ cancels out with $f(x)$ in Equation 6, further reducing the equation to

$$a(x \to x') = \min\{1, \frac{L_{ns,x'}}{L_{ns,x}}\}. \tag{7}$$

$L_{ns,x}$ can be estimated from the first phase of bidirectional sampling for each pixel and hence does not need to be recomputed during the correlated sampling phase.

The reflected radiance at each partially occluded pixel is then estimated as

$$L_{vis}(\omega_r) = \frac{1}{N \cdot C} \sum_{j=1}^{N} \sum_{k=1}^{C} V(\omega_{i,x}^{(j)}) L_{ns,x'} a(\omega_{i,x'}^{(j)} \to \omega_{i,x}^{(j)}), \tag{8}$$

where L_{vis} is the *visibility estimator* in the viewing direction ω_r, $L_{ns,x'} a(\omega_{i,x'}^{(j)} \to \omega_{i,x}^{(j)})$ is the fraction of energy received at pixel x from a neighboring pixel x' during each transition, N is the number of bidirectional samples chosen from first round sampling, and C is the number of energy transitions (Markov chains of length 1) employed in the second round to spread the energy of unoccluded samples, i.e., the valid samples of the target distribution.

4.1 Bias

As we have introduced the method this far, it produces results consistent with the true illumination, but it may exhibit a small amount of bias for finite sample sizes. We use samples for both estimating the direct illumination, and for deciding whether to start Metropolis mutations in partially occluded regions. Such dual use results in a bias, as pointed out by Kirk and Arvo [12], although the bias is typically small – it is less than the standard deviation of the Monte Carlo estimate in the first stage.

In practice, we find this bias to be small enough to be accepted, but if deemed necessary, we can derive a unbiased variant of our method by splitting the MC sample set form the first phase into two partitions: one for deciding whether to apply the Metropolis algorithm, and one used for estimating the illumination in case Metropolis is not necessary. Since we are now using a smaller set of samples to estimate visibility in partially occluded regions, we have to use slightly larger sample chains to achieve the same quality of results. Since the number of visibility tests remains the same, this can be done at low additional cost. Figure 4 shows a comparison of the biased and the unbiased version of the algorithm.

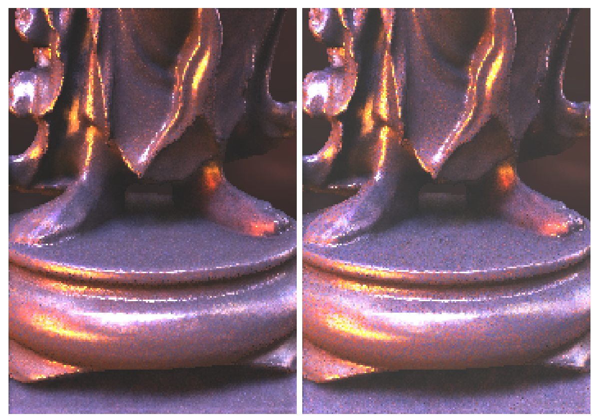

Fig. 4 Visual comparison of the biased (left) and the unbiased version (right) of our method. Note that the highlights are crisper in the unbiased solution.

5 Results

In this section we compare the results of our correlated visibility sampling with bidirectional importance sampling for rendering from HDR environment maps. Images were generated with a reasonably well-optimized ray tracer using a voxel grid as the acceleration data structure for intersection queries. Our comparisons examine the output quality of the two discussed rendering algorithms for a fixed amount of computing time. We performed these tests on a 3.0 GHz Xeon running Linux SuSE 9.0.

Figure 2 presents the visibility masks obtained from first round bidirectional sampling for the images in Figures 1 and 5. The white pixels represent unoccluded pixels which would not be processed in our second round of sampling. The gray pixels correspond to the background environment map. Finally, the black pixels correspond to those where one or more visibility samples were occluded during first round of bidirectional sampling. These pixels are deemed partially occluded and are processed during our second round of correlated visibility sampling.

Fig. 5 Buddha model (Phong BRDF $s = 50$, $k_s = k_d = 0.5$) in an indoor HDR EM. Left: Importance sampling from EM, 200 samples/pixel. Center Left: Multiple importance sampling from EM (140 samples/pixel) and BRDF (40 samples/pixel). Center Right: Bidirectional importance sampling, 20 samples/pixel. Right: Correlated visibility sampling, 16 bidirectional samples (1^{st} pass) and 16 Metropolis samples per unoccluded sample (2^{nd} pass). Rendering times are identical (16 seconds).

Figure 1 presents a complex visibility scenario with the Buddha model in the Grace Cathedral Environment. The Buddha model has 300K triangles, while the Grace Cathedral environment is a 1024×512 HDR map with a contrast ratio of $10^7 : 1$. In this test, both the bidirectional sampling and the correlated sampling algorithms were given 24 seconds to render one 176×248 image each. The time budget was chosen so as to allow good quality in unoccluded regions. For bidirectional sampling, this time budget allowed for visibility to be tested with $N = 20$ samples, and these N samples were chosen after resampling from a larger set of $M = 800$ samples. Note how the shadows between Buddha's feet as well on the ground-plane are noisy with bidirectional sampling. For the same compute time, the partially occluded regions are very smooth with our correlated sampling approach. Here, we used $N = 16$ first round bidirectional samples for the unoccluded regions, and then $C = 16$ Metropolis samples to spread the energy of the unoccluded samples in our second round of sampling. The un-occluded reflected radiance L_{ns} was estimated using fewer samples (M=725) with our correlated sampling approach. Hence, our algorithm produces slightly noisier results in these unoccluded regions. However, the overall tradeoff is better with this approach. The RMS error compared to a converged image reduced from 0.066 when using bidirectional sampling to 0.058 when using the correlated sampling approach for Figure 1. And visually, the images rendered with correlated sampling are much more pleasing due to lower noise in the shadowed regions (please refer the video).

Figure 5 presents a scene with visibility not as complex as that of Figure 1 and with lower frequency illumination. Here we compare the performance of our correlated sampling approach with standard importance sampling from EM, multiple importance sampling from EM and BRDF, as well as bidirectional sampling. Due to high frequencies in both the EM and the BRDF, multiple importance sampling has better performance than sampling only according to the EM. Bidirectional sampling does better than both these approaches in reducing image noise as it samples according to the product distribution. Even then, our correlated sampling approach is more effective than bidirectional sampling in reducing noise in partially occluded regions such as the inside of Buddha's arms and regions around the face and chest.

Figure 7 presents another visibility situation with the Dragon model (870K triangles) in the Grace Cathedral environment. Here, the regions on the Dragon's neck underneath the head as well as on the body are partially occluded by other parts of the Dragon's body. Again, our correlated sampling approach nicely cleans up the shadowed regions that remain noisy with the bidirectional sampling approach. Figure 6 presents the visibility masks for the images in Figure 7. The mask on the top shows many pixels in generally unoccluded areas on the Dragon's body as well as its head are marked as partially occluded after first round of sampling. The mask on the bottom is obtained after applying a simple dilation operation with a 3×3 kernel to the original mask and is better representative of the visibility situation. The dilation operation reduces the variance in the penumbra region, and reproduces sharper shadow boundaries in regions that are generally unoccluded such as the Dragon's head (Figure 7, right).

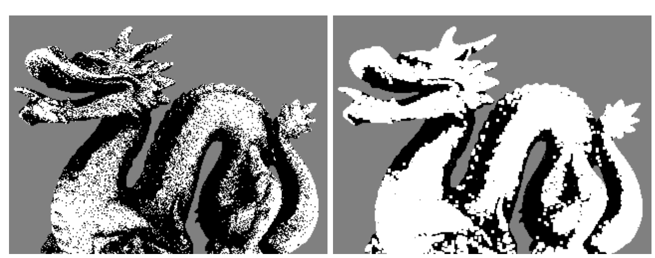

Fig. 6 Visibility mask for the images in Figure 7. Left: Undilated mask. Right: Dilated mask.

Figure 8 presents the David model (700K triangles) with a relatively low frequency BRDF ($s = 10$, $k_s = 0.5$, $k_d = 0.5$) rendered in direct sunlight from an HDR sky probe [28]. No-

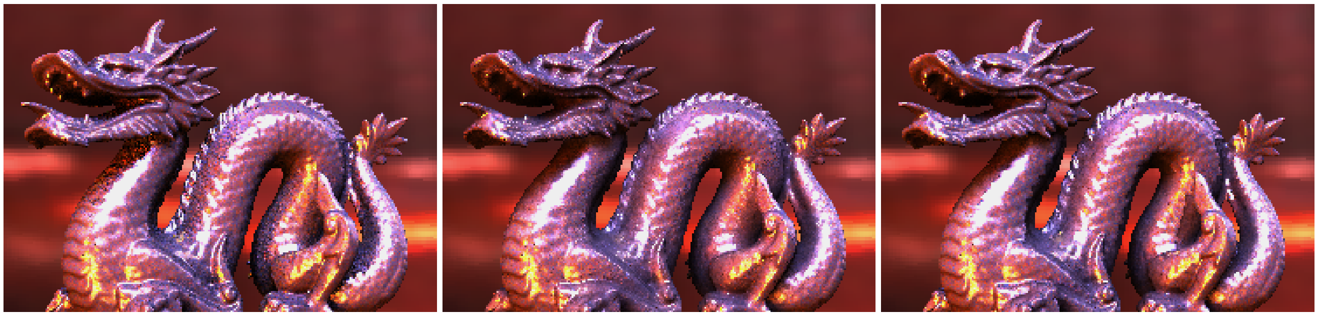

Fig. 7 Dragon model (Phong BRDF $s = 50$, $k_s = k_d = 0.5$) in the Grace Cathedral HDR EM. Left: Bidirectional importance sampling, 20 samples/pixel. Middle and Right: Correlated visibility sampling, 16 bidirectional samples (1^{st} pass) and 16 Metropolis samples per unoccluded sample (2^{nd} pass). Middle: Undilated visibility mask. Right: Dilated visibility mask. Rendering times are identical (16 seconds).

tice how the noise in the shadow on David's chest (Figure 8, left) is efficiently reduced using the correlated sampling approach (Figure 8, right).

Our correlated sampling approach could also be used with traditional importance sampling as the first stage of Monte Carlo sampling instead of bidirectional sampling. However, this would only be efficient when only either the BRDF or the illumination is high frequency but not both, as is the case in Figure 8.

The implementation of our correlated sampling approach involves the usual time vs. memory tradeoff. Compared to the bidirectional sampling approach that processes each pixel independently, the correlated sampling approach needs to store information about neighboring pixels and the visibility mask at the end of the first stage of sampling. For efficiency, we store the N bidirectional samples for each pixel obtained from first stage sampling as well as the estimate of L_{ns} for each pixel. In addition, in order to prevent having to trace primary rays for every transition during correlated sampling, we also store the information corresponding to primary rays such as vertex position, vertex normal and view vector for every pixel. Thus our implementation incurs an additional memory overhead of $\sim W \times H \times N \times Sample$, where $W \times H$ is the resolution of the image plane and $Sample$ is a triple of floats used for storing a sample/position/normal.

With these memory overheads, our correlated sampling stage only required an additional $5-10\%$ computation time after the first stage bidirectional sampling. This additional time was mostly spent in areas with high occlusion such as ground planes occluded by geometry in Figure 1.

Fig. 8 David model (Phong BRDF $s = 10$, $k_s = 0.5, k_d = 0.5$) in a sky probe HDR EM. Left: Bidirectional importance sampling, 20 samples/pixel. Right: Correlated visibility sampling, 16 bidirectional samples (1^{st} pass) and 16 Metropolis samples per unoccluded sample (2^{nd} pass). Rendering times are identical (12 seconds).

6 Conclusions

We have presented a correlated sampling approach for direct illumination that also takes visibility into account in the sampling process apart from the incident illumination and the surface reflectance. By providing a means of sampling from the triple product of the illumination, BRDF and visibility, our method achieves lower variance in partially occluded regions with complex visibility compared to bidirectional sampling for direct illumination. Our method effectively exploits correlation in the integral estimates of neighboring pixels to reduce noise in these regions.

Our proposed correlated sampling method incurs additional memory overheads for storing samples generated from first round of bidirectional sampling. However, this overhead is linear in the number of pixels in the image, and hence is not that significant. Given this small memory overhead, the correlated sampling stage reduces noise in shadowed regions with complex visibility with a small $(5-10\%)$ additional computation time. For this extra computation time, our method of employing Metropolis sampling after an initial phase of Monte Carlo sampling provides much greater benefit in terms of image quality in shadowed regions compared to pure Monte Carlo sampling.

7 Acknowledgments

We would like to thank David Burke for providing us his ray-tracer and Paul Debevec for the HDR sky probe gallery and environment maps. We also thank our anonymous reviewers for their valuable comments and suggestions. The first author was supported by a UBC University Graduate Fellowship for the year 2005-06.

References

1. Agarwal, S., Ramamoorthi, R., Belongie, S., Jensen, H.W.: Structured importance sampling of environment maps. ACM Transactions on Graphics (Proc. SIGGRAPH) **22**(3), 605–612 (2003)
2. Burke, D., Ghosh, A., Heidrich, W.: Bidirectional importance sampling for direct illumination. In: Eurographics Symposium on Rendering, pp. 147–156 (2005)
3. Clarberg, P., Jarosz, W., Akenine-Möller, T., Jensen, H.W.: Wavelet importance sampling: Efficiently evalutating products of complex functions. ACM Transactions on Graphics (Proc. SIGGRAPH) **24**(3), 1166–1175 (2005)
4. Cline, D., Talbot, J., Egbert, P.: Energy redistribution path tracing. ACM Transactions on Graphics (Proc. SIGGRAPH) **24**(3), 1186–1195 (2005)
5. Cohen, J., Debevec, P.: Light-gen. HDRshop plugin (2001). Http://www.ict.usc.edu/j̃cohen/lightgen/lightgen.html
6. Dutre, P., Jensen, H.W.: SIGGRAPH Course Notes: State of the Art in Monte Carlo Global Illumination (2004)
7. Fan, S., Chenney, S., Lai, Y.: Metropolis photon sampling with optional user guidance. In: Eurographics Symposium on Rendering, pp. 127–138 (2005)
8. Greene, N.: Environment mapping and other applications of world projections. IEEE CG&A **6**(11), 21–29 (1986)
9. Heidrich, W., Seidel, H.P.: Realistic, hardware-accelerated shading and lighting. In: Proc. of ACM SIGGRAPH '99, pp. 171–178 (1999)
10. Kautz, J., McCool, M.: Approximation of glossy reflection with prefiltered environment maps. In: Proc. of Graphics Interface, pp. 119–126 (2000)
11. Kautz, J., Vázquez, P.P., Heidrich, W., Seidel, H.P.: Unified approach to prefiltered environment maps. In: Eurographics Workshop on Rendering, pp. 185–196 (2000)
12. Kirk, J., Arvo, J.: Unbiased sampling techniques for image synthesis. In: Proc. of ACM SIGGRAPH '91, pp. 153–156 (1991)
13. Kollig, T., Keller, A.: Efficient illumination by high dynamic range images. In: Eurographics Symposium on Rendering, pp. 45–51 (2003)
14. Lawrence, J., Rusinkiewicz, S., Ramamoorthi, R.: Efficient BRDF importance sampling using a factored representation. ACM Transactions on Graphics (Proc. SIGGRAPH) **23**(3), 496–505 (2004)
15. Lawrence, J., Rusinkiewicz, S., Ramamoorthi, R.: Adaptive numerical cumulative distribution functions for efficient importance sampling. In: Eurographics Symposium on Rendering, pp. 11–20 (2005)
16. Lloyd, S.: An optiumization approach to relaxation labelling algorithms. Image and Vision Computing **1**(2), 85–91 (1983)
17. McCool, M.D., Harwood, P.K.: Probability trees. In: Proc. Graphics Interface, pp. 37–46 (1997)
18. Metropolis, N., Rosenbluth, A.W., Rosenbluth, M.N., Teller, A.H., Teller, E.: Equation of state calculations by fast computing machines. Journal of Chemical Physics **21**(6), 1087–1092 (1953)
19. Ng, R., Ramamoorthi, R., Hanrahan, P.: All-frequency shadows using non-linear wavelet lighting approximation. ACM Transactions on Graphics (Proc. SIGGRAPH) **22**(3), 376–381 (2003)
20. Ostromoukhov, V., Donohue, C., Jodoin, P.M.: Fast hierarchical importance sampling with blue noise properties. ACM Transactions on Graphics (Proc. SIGGRAPH) **23**(3), 488–495 (2004)
21. Ramamoorthi, R., Hanrahan, P.: An efficient representation for irradiance environment maps. In: Proc. of ACM SIGGRAPH '01, pp. 497–500 (2001)
22. Ramamoorthi, R., Hanrahan, P.: Frequency space environment map rendering. In: Proc. of ACM SIGGRAPH '02, pp. 517–526 (2002)
23. Secord, A., Heidrich, W., Streit, L.: Fast primitive distribution for illustration. In: Eurographics Workshop on Rendering, pp. 215–226 (2002)
24. Shirley, P.: Realistic Ray Tracing. A K Peters, Natick (2000)
25. Sloan, P., Kautz, J., Snyder, J.: Precomputed radiance transfer for real-time rendering in dynamic, low-frequency environments. In: Proc. of ACM SIGGRAPH '02, pp. 527–536 (2002)
26. Slusallek, P., Pflaum, T., Seidel, H.P.: Using procedural renderman shaders for global illumination. In: Proc. of Eurographics, vol. 14, pp. 311–324 (1995)
27. Smith, A., Gelfand, A.: Bayesian statistics without tears: a sampling-resampling perspective. pp. 84–88 (1992)
28. Stumpfel, J., Jones, A., Wenger, A., Tchou, C., Hawkins, T., Debevec, P.: Direct hdr capture of the sun and sky (2004). SIGGRAPH 2004 Poster
29. Szecsi, L., Sbert, M., Szirmay-Kalos, L.: Combined correlated and importance sampling in direct light source computation and environment mapping. In: Proc. of Eurographics, vol. 23, pp. 585–593 (2004)
30. Talbot, J., Cline, D., Egbert, P.: Importance resampling for global illumination. In: Eurographics Symposium on Rendering, pp. 139–146 (2005)
31. Veach, E., Guibas, L.: Optimally combining sampling techniques for monte carlo rendering. In: Proc. of ACM SIGGRAPH '95, pp. 419–428 (1995)
32. Veach, E., Guibas, L.: Metropolis light transport. In: Proc. of ACM SIGGRAPH '97, pp. 65–76 (1997)

A. Ghosh Abhijeet Ghosh is currently a PhD candidate in the Department of Computer Science at the University of British Columbia. He received a B.E. in Computer Science & Engineering from Gujarat University, India in 2000, and an M.S. in Computer Science from Stony Brook University, NY in 2002. Ghosh's research interests lie in the field of computer graphics. He is currently interested in realistic rendering, particularly in the development of techniques for HDR acquisition, rendering and display. Ghosh has published several papers on these topics, and has served as external reviewer for leading computer graphics venues such as SIGGRAPH, Eurographics, and Eurographics Symposium on Rendering. He also served on the program committee of Pacific Graphics 2006.

W. Heidrich Wolfgang Heidrich is an Associate Professor in Computer Science at the University of British Columbia. He received a PhD in Computer Science from the University of Erlangen in 1999, and then worked as a Research Associate at the Computer Graphics Group of the Max-Planck-Institute for Computer Science in Saarbrücken, Germany, before joining UBC in 2000. Heidrich's research interests lie in computer graphics, computer vision, and imaging. In particular, he has worked on HDR imaging and displays, sensing and acquisition for graphics, hardware-accelerated and image-based rendering, and global illumination. Heidrich has written over 80 refereed publications on these subjects and has served on numerous program committees. He was the program co-chair for Graphics Hardware 2002, Graphics Interface 2004, and the Eurographics Symposium on Rendering 2006.