Virtual Spherical Gaussian Lights for Real-Time Glossy Indirect Illumination

Yusuke Tokuyoshi
Square Enix Co., Ltd.

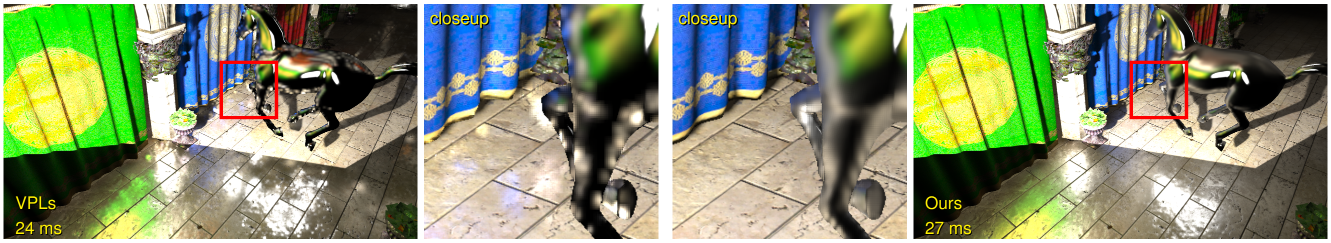

Figure 1: *VPLs (left) and our method (right) for a dynamic scene (264465 triangles). VPLs produce spiky artifacts and temporal flickering, while our method suppresses these artifacts for all-frequency materials with 3 ms overhead (GPU: AMD Radeon™ HD 6990).*

1 Introduction

Virtual point light (VPL) [Keller 1997] based global illumination methods are well established for interactive applications, but they have considerable problems such as spiky artifacts and temporal flickering caused by singularities, high-frequency materials, and discontinuous geometries (Fig. 1). This paper proposes an efficient technique to render one-bounce interreflections for all-frequency materials based on virtual spherical lights (VSLs) [Hašan et al. 2009]. VSLs were proposed to suppress spiky artifacts of VPLs. However, this is unsuitable for real-time applications, since it needs expensive Monte-Carlo (MC) integration and k-nearest neighbor density estimation for each VSL. This paper approximates VSLs using spherical Gaussian (SG) lights without singularities, which take all-frequency materials into account. Instead of k-nearest neighbor density estimation, this paper presents a simple SG lights generation technique using mipmap filtering which alleviates temporal flickering for high-frequency geometries and textures (e.g., normal maps) at real-time frame rates. Since SG lights based approximations are inconsistent estimators, this paper additionally discusses a consistent bias reduction technique. Our technique is simple, easy to integrate in existing reflective shadow map (RSM) based implementations, and completely dynamic for one-bounce indirect illumination including caustics.

2 Background

VPLs and VSLs. For the i-th sampled VPL, reflected radiance at the shading point $\mathbf{x}_s$ and direction $\boldsymbol{\omega}_s$ is given by:

$$L_i(\mathbf{x}_s,\boldsymbol{\omega}_s)=\frac{\rho_i\left(\boldsymbol{\omega}_i\right)}{\|\mathbf{x}_i-\mathbf{x}_s\|^2}, \tag{1}$$

where $\mathbf{x}_i$ is the position of the VPL, and $\boldsymbol{\omega}_i = \frac{\mathbf{x}_i-\mathbf{x}_s}{\|\mathbf{x}_i-\mathbf{x}_s\|}$. The numerator $\rho_i(\boldsymbol{\omega})$ is defined as follows:

$$\rho_i(\boldsymbol{\omega}) = \Phi_i f\left(\mathbf{x}_s,\boldsymbol{\omega}_s,\boldsymbol{\omega}\right) f\left(\mathbf{x}_i,-\boldsymbol{\omega},\boldsymbol{\omega}_i'\right)\langle\boldsymbol{\omega},\mathbf{n}_s\rangle\langle-\boldsymbol{\omega},\mathbf{n}_i\rangle,$$

where Φ_i and $\boldsymbol{\omega}_i'$ are the radiant flux and incident direction of the VPL, $f\left(\mathbf{x}_s,\boldsymbol{\omega}_s,\boldsymbol{\omega}\right)$ is the bidirectional reflectance distribution function (BRDF), $\mathbf{n}_s$ and $\mathbf{n}_i$ are surface normals at $\mathbf{x}_s$ and $\mathbf{x}_i$ respectively, and $\langle\boldsymbol{\omega},\mathbf{n}_s\rangle = \max(\boldsymbol{\omega}\cdot\mathbf{n}_s, 0)$. Although the radiance estimation using Eq. (1) is unbiased and consistent, it can produce high variance, when $\|\mathbf{x}_s-\mathbf{x}_i\|$ is small or BRDFs are high-frequency. In order to suppress this variance, VPLs are approximated with VSLs as follows:

$$L_i(\mathbf{x}_s,\boldsymbol{\omega}_s)\approx\frac{\int_{S^2}\rho_i\left(\boldsymbol{\omega}\right)V_i(\boldsymbol{\omega})d\boldsymbol{\omega}}{\pi r_i^2}. \tag{2}$$

where r_i is the VSL radius estimated by the k-nearest neighbor method (i.e., photon density estimation), and $V_i(\boldsymbol{\omega}) = 1$ if $\boldsymbol{\omega}$ is inside of the cone of directions from $\mathbf{x}_s$ to the VSL, otherwise $V_i(\boldsymbol{\omega}) = 0$. This formulation does not have the singularity caused by $\frac{1}{\|\mathbf{x}_i-\mathbf{x}_s\|^2}$, and the integral can be calculated by MC integration with importance sampling. However, this is unsuitable for real-time applications, since the k-nearest neighbor method and MC integration are expensive.

Analytical approximations. For interactive indirect illumination, analytical area lights approximations have been developed. Prutkin et al. [2012] approximated their VPL clusters with area lights, and used analytical form factors for radiance evaluation. However, this cannot reduce variance caused by highly specular BRDFs. In addition, they did not discuss how to determine the area of their light. Xu et al. [2014] approximated polygon based hierarchical virtual lights (VLs) with SG lights to take all-frequency into account. They approximated the shape of an upper VL node with a disk whose area was the total triangle area. Similar to Xu et al., this paper approximates VSLs using SGs, but VSLs are generated more simply for time-sensitive applications. In addition, this paper estimates the area πr_i^2 based on the variance of VPL positions. This can alleviate flickering artifacts caused by discontinuous geometries for dynamic scenes.

SGs. An SG is a type of spherical function represented by:

$$G(\boldsymbol{\omega},\boldsymbol{\xi},\eta) = e^{\eta((\boldsymbol{\omega}\cdot\boldsymbol{\xi})-1)},$$

where $\boldsymbol{\xi}$ is the lobe axis and η is the lobe sharpness. An anisotropic SG (ASG) [Xu et al. 2013] is defined as:

$$\acute{G}\left(\boldsymbol{\omega},[\boldsymbol{\xi}_x,\boldsymbol{\xi}_y,\boldsymbol{\xi}_z],[\lambda,\mu]\right) = \langle\boldsymbol{\omega},\boldsymbol{\xi}_z\rangle e^{-\lambda(\boldsymbol{\omega}\cdot\boldsymbol{\xi}_x)^2-\mu(\boldsymbol{\omega}\cdot\boldsymbol{\xi}_y)^2},$$

where $\boldsymbol{\xi}_x, \boldsymbol{\xi}_y, \boldsymbol{\xi}_z$ are orthonormal vectors, λ and μ are bandwidth parameters. SGs and ASGs have closed-form solutions of the integral, product, and product integral. Since they are fundamental operators to evaluate illumination integrals, SGs and ASGs are often used for approximating lighting in real-time applications. BRDFs can be represented with an SG-mixture as follows: $f\left(\mathbf{x}_s, \boldsymbol{\omega}_s, \boldsymbol{\omega}\right) \approx \sum_j^K c_j G(\boldsymbol{\omega}, \boldsymbol{\xi}_j, \eta_j)$, where c_j is the lobe coefficient of the SG. For notation simplicity, this paper hereafter rewrites SGs and ASGs as: $G_j(\boldsymbol{\omega}) = G\left(\boldsymbol{\omega}, \boldsymbol{\xi}_j, \eta_j\right)$ and $\acute{G}_j(\boldsymbol{\omega}) = \acute{G}\left(\boldsymbol{\omega}, [\boldsymbol{\xi}_{x,j}, \boldsymbol{\xi}_{y,j}, \boldsymbol{\xi}_{z,j}], [\lambda_j, \mu_j]\right)$. For diffuse-specular models, this can be simplified as: $f\left(\mathbf{x}_s, \boldsymbol{\omega}_s, \boldsymbol{\omega}\right) \approx c_0 + c_1 G_1(\boldsymbol{\omega})$, where c_0 is the coefficient of an SG with zero sharpness to represent the diffuse term. For microfacet BRDFs, Wang et al. [2009] first approximated a normal distribution function (NDF) using SGs, and then approximated the reflection lobe with SGs using *spherical warping* on-the-fly. In this approximation, factors other than the NDF (e.g., the Fresnel term and masking-shadowing function) are represented as c_j. ASGs are also usable in the same manner. Therefore, when incident radiance is represented using SGs or ASGs, the illumination integral can be calculated analytically. Wang et al. [2009] approximated a small sphere light with an SG light for direct illumination. However, since their approximation is not energy conserving, it is inapplicable for approximating a VSL. Xu's SG representation for sphere lights [Xu et al. 2014] is based on energy conservation, however, this can induce the singularity for VSLs since it can have $\frac{1}{\|\mathbf{x}_i - \mathbf{x}_s\|^2}$. This paper represents VSLs using SG lights without this singularity.

3 Our Method

Our rendering framework is based on [Ritschel et al. 2011]. VPLs are generated using bidirectional RSMs (BRSMs), and visibility tests are done by adaptive imperfect shadow maps (AISMs). This paper approximates these generated VPLs with virtual SG lights (VSGLs) on-the-fly. Unlike Xu's polygon based VLs, this easily handles high-frequency textures. In addition, our algorithm also employs this VSGL based radiance evaluation for building the probability density function (PDF) of BRSMs. Therefore, our BRSMs can generate a more accurate sample distribution than the previous BRSMs, especially for high-frequency BRDFs. Since the PDF of BRSMs is computed in the same fashion as VSGLs without visibility tests by exchanging RSMs and a G-buffer, this paper only describes VSGLs generated from RSMs for ease of explanation. Subsect. 3.1 formulates our VSGLs. After that, Subsect. 3.2 describes an efficient VSGL generation and filtering techniques specialized for RSMs. Finally, we discuss the convergence of the VSGL formulation in Subsect. 3.3. This paper assumes BRDFs can be approximated with SGs or ASGs, but they do not have to be extremely accurate for convergence by using a technique described in Subsect. 3.3.

3.1 Virtual Spherical Gaussian Lights

This paper approximates $\int_{S^2} \rho_i\left(\boldsymbol{\omega}\right) V_i(\boldsymbol{\omega}) d\boldsymbol{\omega}$ with a mixture of triple product integrals of two SGs and an ASG. As mentioned in the previous section, BRDFs can be represented using SGs or ASGs to approximate $\rho_i\left(\boldsymbol{\omega}\right)$ on-the-fly. ASGs are more suitable for microfacet BRDFs, because reflection lobes are often anisotropic even if BRDFs have isotropic NDFs. However, the product integral of two ASGs is more computationally expensive and can have a larger precision error than the product integral of an ASG and SG. Moreover, BRDFs on the second bounce are less visually important than the first bounce. Therefore, we employ ASGs and SGs for BRDFs at shading points and BRDFs on VSLs respectively as follows: $f\left(\mathbf{x}_s, \boldsymbol{\omega}_s, \boldsymbol{\omega}\right) \approx \sum_j^K \acute{c}_j \acute{G}_j(\boldsymbol{\omega})$ and $\Phi_i f\left(\mathbf{x}_i, -\boldsymbol{\omega}, \boldsymbol{\omega}'_i\right) \approx \sum_l^K c_l G_l(\boldsymbol{\omega})$. Next, $V_i(\boldsymbol{\omega})$ is represented with an SG using an approximation derived in [Xu et al. 2014] as follows:

$$V_i(\boldsymbol{\omega}) \approx 2G\left(\boldsymbol{\omega}, \boldsymbol{\omega}_i, \frac{4\pi}{\|\Omega_i\|}\right) = 2G_V(\boldsymbol{\omega}),$$

where $\|\Omega_i\|$ is the solid angle of the cone, and $\|\Omega_i\| \approx \frac{\pi r_i^2}{\|\mathbf{x}_i - \mathbf{x}_s\|^2}$ is used in [Xu et al. 2014]. However, this $\|\Omega_i\|$ cannot avoid the singularity caused by $\frac{1}{\|\mathbf{x}_i - \mathbf{x}_s\|^2}$ for VSLs when BRDFs are diffuse models. To avoid this singularity, this paper employs $\|\Omega_i\| = \int_{S^2} V_i(\boldsymbol{\omega}) d\boldsymbol{\omega} = 2\pi(1 - \cos\theta_i)$ similar to the original VSL method, where θ_i is the angle of the cone given by $\cos\theta_i = \frac{\|\mathbf{x}_i - \mathbf{x}_s\|}{\sqrt{\|\mathbf{x}_i - \mathbf{x}_s\|^2 + r_i^2}}$. Finally, since lobe coefficients and cosine terms are smooth and do not induce spiky artifacts, they are pulled out of the integral in Eq. (2). Thus, we obtain the following equation:

$$\int_{S^2} \rho_i\left(\boldsymbol{\omega}\right) V_i(\boldsymbol{\omega}) d\boldsymbol{\omega} \approx \sum_j^K \sum_l^K q_{j,l} \int_{S^2} \acute{G}_j(\boldsymbol{\omega}) G_l(\boldsymbol{\omega}) G_V(\boldsymbol{\omega}) d\boldsymbol{\omega}, \tag{3}$$

where $q_{j,l} = 2\acute{c}_j c_l \langle \boldsymbol{\omega}_i, \mathbf{n}_s \rangle \langle -\boldsymbol{\omega}_i, \mathbf{n}_i \rangle$. We calculate this triple product integral analytically instead of using MC integration. The accuracy of the BRDF representation can be improved by increasing K, but this is inefficient because Eq. (3) has complexity $O\left(K^2\right)$. Therefore, this paper restricts $K = 2$ for diffuse-specular BRDFs. Since this representation has zero sharpness for a diffuse term, Eq. (3) can be computed more simply. Even if multi-lobe BRDFs are approximated using small K, this approximation error can converge to zero for an infinite number of VSGLs using a technique described in Sect. 3.3.

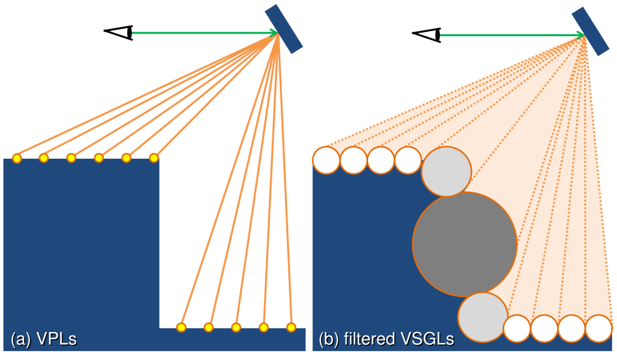

Figure 2: *VPLs (a) and our filtered VSGLs (b). Due to filtering, larger and darker VSGLs are generated on discontinuous geometries. This approach produces smoother results, and thus alleviates temporal flickering.*

3.2 Filtered Importance Sampling for RSMs

VSGLs can be generated in the same way as VSLs, but the k-nearest neighbor method is expensive for real-time applications. In addition, flickering caused by high-frequency textures and discontinuous geometries are a considerable issue. To avoid these problems for one-bounce indirect illumination, this paper generates VSGLs from RSMs by using filtered importance sampling [Křivánek and Colbert 2008]. This technique smoothes sampled radiance on unimportant regions, and it is often used for variance reduction of environmental lighting. We employ it for RSMs to approximate light distribution. In our algorithm, sample points are selected in the same fashion as BRSM technique, but their values are obtained from the mipmapped RSMs. The sampled miplevel h is determined based on the sample density as follows:

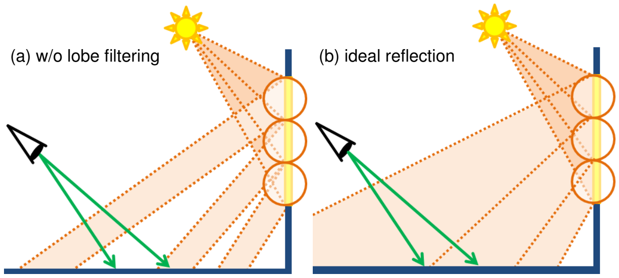

Figure 3: *For caustics, VSGLs without lobe filtering can produce spot artifacts (a). To approximate the ideal reflection (b), this paper employs SG lobe mipmap filtering.*

$h = h_{max} - \frac{1}{2}\log_2\left(\frac{N}{k}p\left(\boldsymbol{\omega_i}'\right)\right)$, where h_{max} is the top miplevel of the RSM, N is the number of VSGLs, $p(\boldsymbol{\omega_i}')$ is the PDF of BRSMs, and k is the user-specified parameter to control the smoothness similar to the k-nearest neighbor method. Using this simple technique, $\mathbf{x}_i$ and $\mathbf{n}_i$ can be filtered to alleviate flickering artifacts.

Radius estimation. The radius r_i should be estimated based on the positional distribution of VPLs in world space similar to VSLs. Therefore, this paper estimates r_i using the positional variance which is also calculated using the filtered values. This positional variance is given as: $\sigma^2 = y_i - \|\mathbf{x}_i\|^2$, where y_i and $\mathbf{x}_i$ are the sampled filtered values of $\|\mathbf{x}\|^2$ and $\mathbf{x}$ stored in RSMs respectively, and $\mathbf{x}$ is the position at each RSM texel. Assuming a uniform distribution on a disk, r_i is given by: $r_i^2 = 2\sigma^2$. For discontinuous geometries, although this produces larger radii which induce smoother results, it can be more visually acceptable for dynamic scenes since flickering artifacts are alleviated.

Filtering SG lobes. This paper employs $K = 2$, and approximates diffuse BRDFs with an SG with zero sharpness. Therefore, diffuse coefficients are stored in RSMs, and then simply filtered by mipmapping. For specular lobes, filtering techniques for directional distribution functions such as [Toksvig 2005; Laurijssen et al. 2010] can be used to suppress spot artifacts for caustics as illustrated in Fig. 3. Since these techniques can represent a normalized SG with its averaged direction, a filtered SG lobe is obtained by averaging operations such as mipmapping. They are also usable for reducing K on-the-fly for multi-lobe BRDFs. In the RSM rendering pass, a specular BRDF is approximated with a normalized SG and its coefficient, and then this normalized SG is represented with an averaged direction. This direction and coefficient are stored in each RSM texel. For VSGL generation, the filtered SG lobe is restored from the mipmapped averaged direction and coefficient. Although Laurijssen's approximation is more accurate than Toksvig's approximation for SGs, it has to solve a cubic equation to convert an SG into an averaged direction. On the other hand, Toksvig's approximation solves a linear equation. Our current implementation employs Toksvig's approximation for simplicity.

To perform the above techniques our RSM has positions, squared norms of the positions, normals, diffuse coefficients, specular coefficients, and averaged directions of specular lobes. This simple approach generates VSGLs at real-time frame rates with less noticeable artifacts.

3.3 Consistent Bias Reduction

Our method produces visually acceptable results by exchanging variance for bias. However, bias produced by spherical warping cannot converge to zero, even if an NDF is well approximated using $K \to \infty$. Thus, SG light based methods using such approximations are biased and inconsistent estimators. However, for our method, this bias can converge to zero with a slight modification. This is done by canceling the SG approximation of BRDFs while increasing the number of VSGLs. This cancellation is determined by comparing $\rho_i(\boldsymbol{\omega})$ and $V_i(\boldsymbol{\omega})$ in the SG lobe sharpness metric. Since $\rho_i(\boldsymbol{\omega}) \approx \sum_j^K \sum_l^K \frac{q_{j,l}}{2}\acute{G}_j(\boldsymbol{\omega})G_l(\boldsymbol{\omega})$ which is closed in ASG basis, its maximum sharpness η_ρ can be computed analytically. In addition, the sharpness of $V_i(\boldsymbol{\omega})$ is $\frac{4\pi}{\|\Omega_i\|}$. Therefore, the comparison of $\rho_i(\boldsymbol{\omega})$ and $V_i(\boldsymbol{\omega})$ can be represented as: $\eta_\rho < \alpha\frac{4\pi}{\|\Omega_i\|}$, where α is the user-specified parameter. If this condition is true (i.e., $\rho_i(\boldsymbol{\omega})$ is relatively smooth), $\rho_i(\boldsymbol{\omega})$ can be pulled out of the integral of Eq. (2) as follows:

$$L_i(\mathbf{x}_s, \boldsymbol{\omega}_s) \approx \frac{\rho_i\left(\boldsymbol{\omega}_i\right)\int_{S^2} V_i(\boldsymbol{\omega})d\boldsymbol{\omega}}{\pi r_i^2} = \frac{\rho_i\left(\boldsymbol{\omega}_i\right)\|\Omega_i\|}{\pi r_i^2}.$$

Increasing the RSM resolution and number of VSGLs, the filtering kernel width, r_i, and $\|\Omega_i\|$ become smaller. In addition, for $r_i \to 0$, this equation converges to Eq. (1). Hence, our formulation can be theoretically consistent regardless of K.

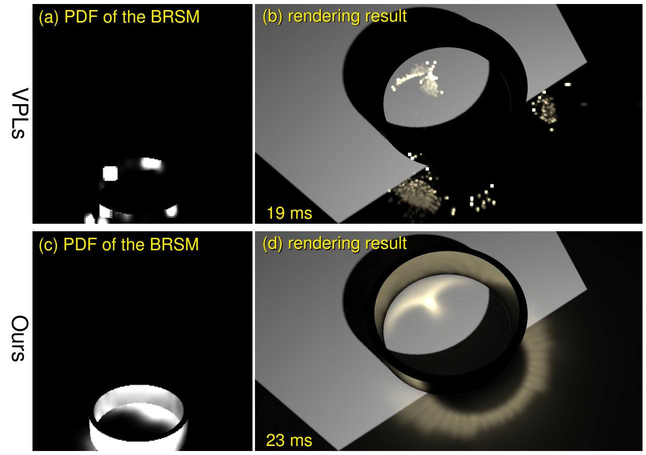

Figure 4: *Caustics (514 triangles, Phong exponent: 65535). The original BRSM technique cannot produce an appropriate PDF for specular surfaces (a), and thus generates VPLs with high variance. In addition, the resulting image has spiky artifacts (b). Our method can render more accurate results (c)(d) without such artifacts.*

4 Experimental Results

In the following, we show the rendering results using $N = 1024$ and $k = 4$ performed on an AMD Radeon HD 6990. These experimental scenes have microfacet BRDFs with Phong NDFs. The frame buffer and RSM resolution are 1920×1088 and 256^2, respectively. For each AISM, 8192 points and 64^2 resolution are used.

Quality. Fig. 4 shows caustics rendered using VPLs (upper) and our VSGLs (lower) at real-time frame rates. The original BRSM technique (a) is unsuitable for such scenes since the PDF can be estimated with large variance which induces additional errors for resulting images. In addition, the VPL evaluation also has large variance. Hence, intense spiky artifacts and temporal flickering are produced (b). On the other hand, our method renders such caustics with less variance (c)(d). Fig. 5 shows a textured scene with complex geometries and normal maps. For this scene, 1024 VPLs

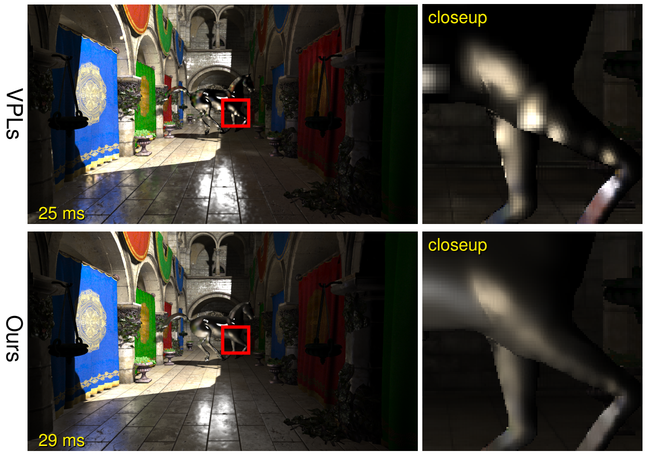

Figure 5: *Textured scene with normal maps (The same scene as Fig. 1, Phong exponent: 1023). Using filtered VSGLs, spiky artifacts and temporal flickering are reduced significantly.*

Table 1: *Computation times for indirect illumination (ms).*

	Fig. 4		Fig. 5	
	VPLs	Ours	VPLs	Ours
G-buffer	0.51	0.51	1.54	1.54
RSMs	0.06	0.06	0.41	0.41
PDF of BRSMs	**2.45**	**2.78**	**2.76**	**3.05**
VPL/VSGL generation	**0.04**	**0.35**	**0.04**	**0.38**
AISMs	7.65	8.76	12.20	12.39
Radiance estimation	**7.06**	**9.63**	**6.93**	**9.56**

are insufficient to accurately represent indirect illumination. Our method significantly reduces spiky artifacts and temporal flickering by filtering. For such insufficient samples, although our approach can produce lower-frequency illumination appearance than is actually the case, it is visually acceptable compared to the VPL based method which produces high-frequency artifacts.

Performance. Table 1 shows the computation time of each procedure in Figs. 4 and 5. Our contributions are written in bold. These times are almost scene independent. The difference of AISMs is caused by the error of the PDF. The overhead of VSGL generation is due to mipmapping of RSMs. The total overhead is 3-4 ms.

Bias reduction. Rendering results using the consistent bias reduction technique are shown in Fig. 6. While this can produce more visually acceptable results for small N, the bias can converge to zero with increasing N by canceling the SG approximation without noticeable spiky artifacts. However, our current implementation uses AISMs which induce additional bias and temporal flickering. The avoidance of the error caused by AISMs is our future work.

5 Conclusion and Future Work

This paper has presented a simple approximation of VSLs called VSGLs using SG based radiance evaluation and filtered importance sampling specialized for RSMs. Using this technique, we are able to render one-bounce glossy interreflections including caustics with a few milliseconds overhead compared to VPLs. When the number of VPLs is insufficient to approximate indirect illumination, the proposed method exchanges undesirable variance of such VPLs for visually acceptable bias in a simplistic way. Furthermore, this paper has discussed the modification for consistent estimation. Although our current implementation uses AISMs to perform real-time rendering, our estimator is also potentially applicable to ray tracing based renderers. In addition, since VSL radius estimation is based on photon density estimation, our algorithm can be extended to progressive rendering by multiplying the radii by $\frac{2}{3}$ for each iteration, similar to progressive photon mapping [Hachisuka et al. 2008]. We would like to investigate the efficiency of such algorithms for offline or near-interactive rendering in the future.

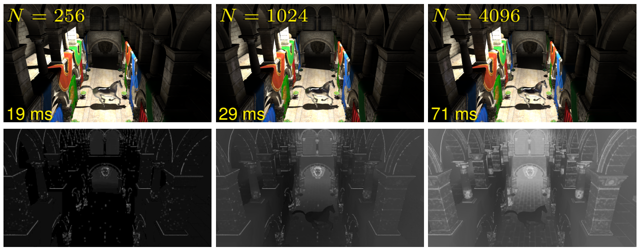

Figure 6: *Canceling the SG approximation of BRDFs using* $\alpha = 0.001$ *for the same scene as Fig. 1. Upper: rendered images. Lower: visualization of the cancellation rate. The SG approximation is canceled with increasing the number of VSGLs, and thus the bias can converges to zero without noticeable artifacts.*

References

HACHISUKA, T., OGAKI, S., AND JENSEN, H. W. 2008. Progressive photon mapping. *ACM Trans. Graph. 27*, 5, 130:1–130:8.

HAŠAN, M., KŘIVÁNEK, J., WALTER, B., AND BALA, K. 2009. Virtual spherical lights for many-light rendering of glossy scenes. *ACM Trans. Graph. 28*, 5, 143:1–143:6.

KELLER, A. 1997. Instant radiosity. In *Proc. SIGGRAPH'97*, 49–56.

KŘIVÁNEK, J., AND COLBERT, M. 2008. Real-time shading with filtered importance sampling. *Comput. Graph. Forum 27*, 4, 1147–1154.

LAURIJSSEN, J., WANG, R., DUTRÉ, P., AND BROWN, B. J. 2010. Fast estimation and rendering of indirect highlights. *Comput. Graph. Forum 29*, 4, 1305–1313.

PRUTKIN, R., KAPLANYAN, A. S., AND DACHSBACHER, C. 2012. Reflective shadow map clustering for real-time global illumination. In *Eurographics 2012 Short Papers*, 9–12.

RITSCHEL, T., EISEMANN, E., HA, I., KIM, J. D., AND SEIDEL, H.-P. 2011. Making imperfect shadow maps view-adaptive: High-quality global illumination in large dynamic scenes. *Comput. Graph. Forum 30*, 8, 2258–2269.

TOKSVIG, M. 2005. Mipmapping normal maps. *J. Graph. Tools 10*, 3, 65–71.

WANG, J., REN, P., GONG, M., SNYDER, J., AND GUO, B. 2009. All-frequency rendering of dynamic, spatially-varying reflectance. *ACM Trans. Graph. 28*, 5, 133:1–133:10.

XU, K., SUN, W.-L., DONG, Z., ZHAO, D.-Y., WU, R.-D., AND HU, S.-M. 2013. Anisotropic spherical gaussians. *ACM Trans. Graph. 32*, 6, 209:1–209:11.

XU, K., CAO, Y.-P., MA, L.-Q., DONG, Z., WANG, R., AND HU, S.-M. 2014. A practical algorithm for rendering interreflections with all-frequency brdfs. *ACM Trans. Graph. 33*, 1, 10:1–10:16.