

# Stratified Sampling of Projected Spherical Caps

Carlos Ureña<sup>1</sup>

Iliyan Georgiev<sup>2</sup>

<sup>1</sup>Universidad de Granada

<sup>2</sup>Solid Angle

---

## Abstract

*We present a method for uniformly sampling points inside the projection of a spherical cap onto a plane through the sphere's center. To achieve this, we devise two novel area-preserving mappings from the unit square to this projection, which is often an ellipse but generally has a more complex shape. Our maps allow for low-variance rendering of direct illumination from finite and infinite (e.g. sun-like) spherical light sources by sampling their projected solid angle in a stratified manner. We discuss the practical implementation of our maps and show significant quality improvement over traditional uniform spherical cap sampling in a production renderer.*

## CCS Concepts

•Computing methodologies → Rendering; Ray tracing; Visibility;

---

## 1. Introduction

A key ingredient in achieving visual realism in computer-generated images is the use of soft illumination from area light sources. Spherical luminaires are commonly utilized in modern production rendering as a substitute for classical point lights to provide both smooth lighting and soft shadows. When placed at infinity, such lights can also simulate directional illumination from the sun.

Monte Carlo integration is the standard method for estimating the illumination from an area light source at a point [SWZ96, PJH16]. This involves sampling the radiance contribution from directions inside the solid angle subtended by the luminaire. Doing that accurately is crucial for minimizing the amount of noise in the rendered image. A good strategy is to distribute the directions uniformly inside the subtended solid angle by transforming stratified canonical unit-square sample points. This requires devising a specialized mapping for each type of light source. Spherical lights have been among the first to be addressed as their solid angle is a simple spherical cap [Wan92]. To facilitate stratification and improve sampling quality, area-preserving solid angle maps have been developed for triangles [Arv95] and rectangles [UFK13]. Gamito [Gam16] proposed a method for cylindrical and disk luminaires based on bounding their solid angle by simpler shapes followed by rejection sampling. Recently, Guillén et al. [GUnK\*17] derived analytical parameterizations of the solid angle of disks.

When the shading point is on a surface, an even more efficient strategy is to importance sample the cosine illumination foreshortening term. According to the Nusselt analog [CWH93, p. 80], doing this is equivalent to sampling points uniformly inside the projection of the solid angle on the shading point's tangent plane. Unfortunately, this is hard to achieve for an arbitrary-shaped luminaire as

it involves first computing and then uniformly sampling this projection. Even if the solid angle has a simple shape, e.g. a spherical cap for a spherical light, the shape of its projection onto the tangent plane is not necessarily simple too. For a spherical triangle, Ureña [Un00] proposed to adaptively subdivide its projection into smaller triangles with nearly constant cosine and then uniformly sample their solid angle [Arv95]. Recent work on spherical cap preserving transformations on the sphere [DHB17] provides an approximate method for producing samples distributed according to projected area (clamped cosine distribution) in an arbitrary spherical cap. To our knowledge, the only existing stratified exact solid angle sampling technique is that of Arvo [Arv01] for polygonal light sources. Sampling the projection of a solid angle is generally more computationally expensive than sampling the solid angle. However, in most practical scenes, where the sample contribution evaluation is orders of magnitude more costly than the sample generation, the resulting lower variance improves the overall efficiency.

In this paper, we present practical techniques for projected solid angle sampling of spherical light sources. To this end, we devise two analytical area-preserving maps from the unit square to the projection of a spherical cap at a surface point (see Figure 1). Our maps can be used to warp stratified points on the unit square into stratified points on that projection, whose shape is often an ellipse but is generally more complex. Besides allowing for sampling only the visible parts of the sphere, our maps are continuous under continuous change in the sphere's position relative to the shading point, which prevents noise discontinuities in the rendered image. We discuss the efficient practical implementation of our method and demonstrate the efficiency improvement it brings over uniform spherical cap sampling in the Arnold production renderer.

## 2. Problem statement

Our goal is to compute the outgoing radiance at a surface point  $\mathbf{x}$  in direction  $\omega_0$  due to scattered direct illumination from a sphere-shaped light source  $\mathcal{S}$ . This radiance is given by an integral over the spherical cap  $\mathcal{C}$  subtended by the light source at  $\mathbf{x}$ :

$$L_o(\mathbf{x}, \omega_0) = \int_{\mathcal{C}} L_i(\mathbf{x}, \omega) f_s(\mathbf{x}, \omega_0, \omega) |\cos \theta| d\sigma(\omega), \quad (1)$$

where  $L_i(\mathbf{x}, \omega)$  is the incident radiance at  $\mathbf{x}$  from direction  $\omega$ , and  $f_s$  is the bidirectional scattering distribution function (BSDF) describing the scattering properties at  $\mathbf{x}$ . The term  $|\cos \theta|$  accounts for the light foreshortening effect at grazing angles  $\theta$  between  $\omega$  and the surface normal at  $\mathbf{x}$ , and  $d\sigma$  is the differential solid angle measure.

Monte Carlo estimation of Equation (1) using  $N$  sample directions  $\omega_i$  has the following form:

$$L_o(\mathbf{x}, \omega_0) \approx \frac{1}{N} \sum_{i=1}^N \frac{L_i(\mathbf{x}, \omega_i) f_s(\mathbf{x}, \omega_i, \omega_0) |\cos \theta|}{p(\omega_i)}. \quad (2)$$

Traditionally, *uniform solid angle sampling* is used to generate directions with a constant pdf  $p(\omega_i) = 1/\sigma(\mathcal{C})$  inside the spherical cap  $\mathcal{C}$ . This technique is easy to implement [Wan92, SWZ96] but can cause high variance in the radiance estimator (2), especially when the light source illuminates the point  $\mathbf{x}$  from a grazing angle.

### 2.1. Stratified projected solid angle sampling

While the incident radiance  $L_i$  and BSDF  $f_s$  are scene- and position-dependent, the cosine foreshortening is always the same. It is therefore sensible to importance sample this term when estimating Equation (1), i.e. to factor it in the sample distribution. To this end, we first express the direction  $\omega$  as a function of its projection  $\mathbf{w}$  on the tangent plane at  $\mathbf{x}$ . It holds that  $dA(\mathbf{w})/d\sigma(\omega) = |\cos \theta|$ , where  $dA$  is the differential area measure. We then make a change of variable in Equation (1) to write it as an integral over the projection  $\mathcal{C}^\perp$  of the spherical cap:

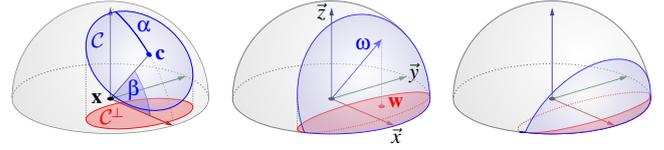
$$\begin{aligned} L_o(\mathbf{x}, \omega_0) &= \int_{\mathcal{C}^\perp} L_i(\mathbf{x}, \omega(\mathbf{w})) f_s(\mathbf{x}, \omega_0, \omega(\mathbf{w})) dA(\mathbf{w}) \\ &= \int_{[0,1]^2} L_i(\mathbf{x}, \omega(\mathbf{w}(s,t))) f_s(\mathbf{x}, \omega_0, \omega(\mathbf{w}(s,t))) J_{\mathbf{w}}(s,t) ds dt. \end{aligned} \quad (3)$$

In Equation (4) we make an additional change of variable to express point  $\mathbf{w} \in \mathcal{C}^\perp$  as a function of a point  $(s,t)$  on the unit square  $[0,1]^2$ .

In this paper we seek to devise mappings from the unit square to the projected spherical cap  $\mathcal{C}^\perp$  that have a constant Jacobian  $J_{\mathbf{w}}(s,t) = A(\mathcal{C}^\perp)$ . This property preserves relative areas from one domain to the other, which allows for easy stratification by using jittered or low-discrepancy unit-square patterns. The estimator for  $L_o$  using  $N$  uniformly distributed unit-square samples then reads

$$L_o(\mathbf{x}, \omega_0) \approx \frac{A(\mathcal{C}^\perp)}{N} \sum_{i=1}^N L_i(\mathbf{x}, \omega(\mathbf{w}(s_i, t_i))) f_s(\mathbf{x}, \omega_0, \omega(\mathbf{w}(s_i, t_i))). \quad (5)$$

Note that this strategy corresponds to importance sampling the cosine term in estimator (2) by using a pdf  $p(\omega) = |\cos \theta|/A(\mathcal{C}^\perp)$ . Note also that when  $L_i$  and  $f_s$  are constant, i.e. the light source emission is uniform and unoccluded and the BSDF is diffuse, estimator (5) always returns the exact result, even with a single sample.



**Figure 1:** A spherical cap  $\mathcal{C}$  (blue) and its projection  $\mathcal{C}^\perp$  (red) on the tangent plane, parameterized in a reference frame aligned with the surface normal and the cap center. The shape of  $\mathcal{C}^\perp$  can be an ellipse (left), the union of an ellipse and a lune (middle), or just a lune (right), respectively when  $\mathcal{C}$  is entirely above the horizon, mostly above the horizon, or mostly below the horizon.

### 2.2. Overview

The radiance estimator (5) uses the Nusselt analog to reduce the task of sampling cosine-distributed directions  $\omega$  on the spherical cap  $\mathcal{C}$  to the task of sampling uniform points  $\mathbf{w}$  on its tangent-plane projection  $\mathcal{C}^\perp$  (see Figure 1, middle). This can in turn be achieved by using an area-preserving mapping to transform stratified unit-square sample points to stratified points in  $\mathcal{C}^\perp$ . We devise two such maps that differ in the way they parameterize the area of  $\mathcal{C}^\perp$ .

Given a point  $(s,t) \in [0,1]^2$ , the idea of both our maps is to first find a line segment that cuts a portion of  $\mathcal{C}^\perp$  with area  $tA(\mathcal{C}^\perp)$  and then to sample a distance along that segment using  $s$ . Our *parallel map* cuts slices of  $\mathcal{C}^\perp$  enclosed between two parallel lines, and our *radial map* cuts sectors of  $\mathcal{C}^\perp$  by sweeping a radial line about its center.

The remainder of the paper is organized as follows. In the following Section 3 we compute the geometry of the spherical cap  $\mathcal{C}$  and its projection  $\mathcal{C}^\perp$ . We then present our area-preserving maps in Section 4. In Section 5, we discuss the practical implementation of our maps and compare their sampling qualities against traditional uniform spherical cap sampling. We draw final conclusions and discuss future work in Section 6.

## 3. Geometric configuration

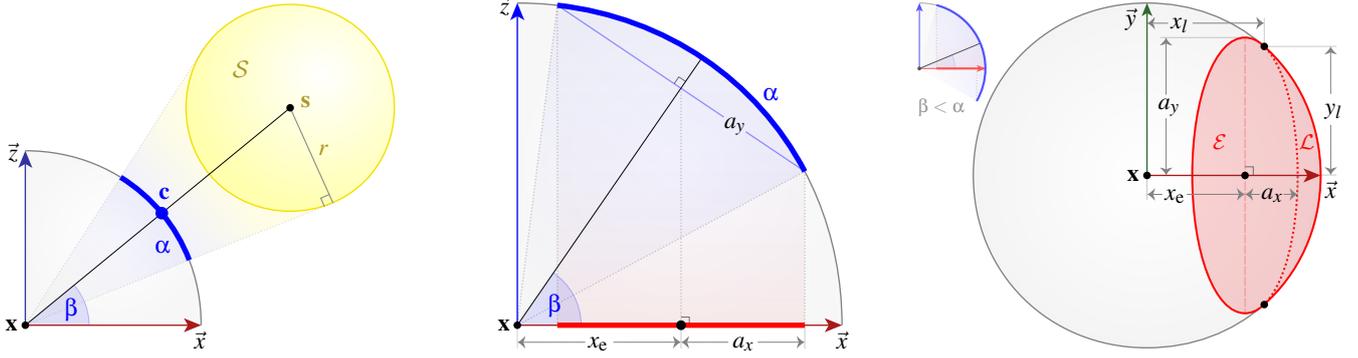
In this section we take the first step toward deriving our maps, which is to choose a reference frame and parameterize the geometry of the spherical cap  $\mathcal{C}$  and its projection  $\mathcal{C}^\perp$  in it. These are illustrated in Figures 1 and 2.

### 3.1. Spherical cap

Given a surface shading point  $\mathbf{x}$  with normal vector  $n_{\mathbf{x}}$  and a spherical cap  $\mathcal{C}$  with center  $\mathbf{c}$  and aperture angle  $\alpha$ , we carry out our derivations in a local reference frame  $\mathcal{R} = (\vec{x}, \vec{y}, \vec{z})$  centered at  $\mathbf{x}$ :<sup>†</sup>

$$\vec{z} = n_{\mathbf{x}}, \quad \vec{y} = \frac{\vec{z} \times (\mathbf{c} - \mathbf{x})}{\|\vec{z} \times (\mathbf{c} - \mathbf{x})\|}, \quad \vec{x} = \vec{y} \times \vec{z}. \quad (6)$$

<sup>†</sup> When the light source is right above  $\mathbf{x}$ , i.e.  $n_{\mathbf{x}}$  is (almost) parallel to  $\mathbf{c} - \mathbf{x}$ , we can choose  $\vec{y}$  to be any vector perpendicular to the normal. This however may yield incoherent reference frames between nearby shading points  $\mathbf{x}$ .



**Figure 2:** Parameterizations of the spherical cap  $C$  and its tangent-plane projection  $C^\perp$  whose shape can be an ellipse  $\mathcal{E}$  and/or a lune  $\mathcal{L}$ . Left: In our reference frame centered at the shading point  $\mathbf{x}$ , the cap is fully specified by two angles: the aperture  $\alpha$  and the elevation  $\beta$ . Middle: The ellipse is specified by its location  $x_e$  along the  $\bar{x}$ -axis and its semi-minor and semi-major axes  $a_x$  and  $a_y$ , computed from  $\alpha$  and  $\beta$ . Right: When  $\beta < \alpha$ , a lune is formed where the ellipse touches the unit circle at distances  $x_l$  and  $y_l$  along the  $\bar{x}$ -axis and  $\bar{y}$ -axis, respectively.

A point  $\mathbf{w} = (x, y)$  on the  $xy$ -plane can be projected up onto the sphere to obtain a world-space direction  $\omega$  (see Figure 1, middle):

$$\omega(\mathbf{w}) = \bar{x}x + \bar{y}y + \bar{z}\sqrt{1-x^2-y^2}. \quad (7)$$

Infinitely distant spherical light sources, e.g. the sun, are readily specified as spherical caps with center  $\mathbf{c}$  and aperture angle  $\alpha$ . In our reference frame it is convenient to use the elevation angle  $\beta$  in lieu of  $\mathbf{c}$  (see Figure 1, left, and Figure 2). For a finite spherical source  $S$  with center  $\mathbf{s}$  and radius  $r$ , these three parameters are:

$$\mathbf{c} = \frac{\mathbf{s} - \mathbf{x}}{\|\mathbf{s} - \mathbf{x}\|}, \quad \alpha = \arcsin\left(\frac{r}{\|\mathbf{s} - \mathbf{x}\|}\right), \quad \beta = \arcsin(\bar{z} \cdot (\mathbf{c} - \mathbf{x})). \quad (8)$$

The cap  $C$  is fully parameterized by the aperture  $\alpha \in [0, \pi/2]$  and elevation  $\beta \in [-\pi/2, \pi/2]$ .<sup>‡</sup> Figure 2, left, illustrates all quantities.

### 3.2. Projected spherical cap

The perimeter of the spherical cap  $C$  is a circle whose parallel projection on the  $xy$ -plane is an ellipse  $\mathcal{E}$ . When  $\alpha > |\beta|$ ,  $C$  intersects the  $xy$ -plane and  $\mathcal{E}$  touches the unit circle (see Figure 1) at two locations. A lune  $\mathcal{L}$  is thus formed between the ellipse and the unit circle (see Figure 2, right). For the projection  $C^\perp$  of the spherical cap we distinguish between four cases:

- $C$  is entirely above the horizon:  $0 \leq \alpha \leq \beta$ . Then  $C^\perp \equiv \mathcal{E}$ .
- $C$  is mostly above the horizon:  $0 \leq \beta < \alpha$ . Then  $C^\perp \equiv \mathcal{E} \cup \mathcal{L}$ .
- $C$  is mostly below the horizon:  $0 < -\beta < \alpha$ . Then  $C^\perp \equiv \mathcal{L}$ .
- $C$  is entirely below the horizon:  $0 \leq \alpha \leq -\beta$ . Then  $C^\perp \equiv \emptyset$ .

If the surface can receive illumination from both the upper and the lower hemispheres around  $\mathbf{x}$ , one of (a)-(c) cases applies for each, and the two hemispheres can be handled separately.

<sup>‡</sup> Note that the slightly more intuitive definition of the elevation angle  $\beta = \arccos(\bar{x} \cdot (\mathbf{c} - \mathbf{x}))$  would give an incorrect range  $\beta \in [0, \pi/2]$ .

**Ellipse.** The ellipse  $\mathcal{E}$  is parameterized by its location  $x_e$  along the  $\bar{x}$ -axis and its semi-minor and semi-major axes  $a_x$  and  $a_y$  (see Figure 2). Since the projection of the circular spherical cap perimeter onto the  $xy$ -plane preserves distances along the  $\bar{y}$ -axis,  $a_y$  is equal to the circle radius. The expressions for these three parameters then follow via basic trigonometry from Figure 2, middle:

$$a_y = \sin \alpha, \quad a_x = a_y \sin \beta = \sin \alpha \sin \beta, \quad x_e = \cos \alpha \cos \beta. \quad (9)$$

Note that when the spherical cap center goes below the horizon, the ellipse flips about its semi-major axis, making  $\beta < 0$  and  $a_x < 0$ .

**Lune.** The lune  $\mathcal{L}$  is parameterized by the locations of the ellipse-circle tangency points with coordinates  $(x_l, y_l)$  and  $(x_l, -y_l)$ , as shown in Figure 2, right. Since the tangency points lie on both the ellipse and the unit circle, from the ellipse and circle equations,

$$E(x_l - x_e, y_l) = \frac{(x_l - x_e)^2}{a_x^2} + \frac{y_l^2}{a_y^2} = 1, \quad C(x_l, y_l) = x_l^2 + y_l^2 = 1, \quad (10)$$

we can obtain simple expressions for their coordinates:

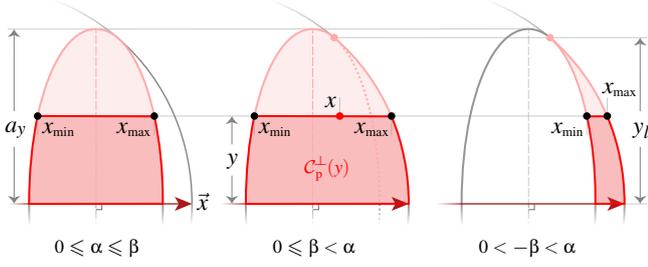
$$x_l = \frac{a_y^2 x_e}{a_y^2 - a_x^2} = \frac{\cos \alpha}{\cos \beta}, \quad y_l = \sqrt{1 - x_l^2} = \sqrt{1 - \frac{\cos^2 \alpha}{\cos^2 \beta}}. \quad (11)$$

Note that  $\mathcal{L}$  is defined and has non-zero area only when  $|\beta| < \alpha$ .

### 4. Area-preserving maps

Having characterized the geometry of the spherical cap projection  $C^\perp$ , we can now derive our two maps for transforming points on the unit-square  $[0, 1]^2$  into points inside  $C^\perp$ .

The idea of our both maps is to first find a line that cuts a portion of  $C^\perp$  with a desired area and then to sample a point along that line using a simple analytical distribution. The challenging part is the derivation of an expression for the area of the cut portion as a function of the line location, and then finding the inverse of that expression. We derive analytical expressions for the portion's area, which we invert numerically. Since the shape of  $C^\perp$  is symmetric about the  $\bar{x}$ -axis (see Figure 2, right), we consider only the half that lies on the positive side of the  $\bar{y}$ -axis.



**Figure 3:** Our parallel map cuts slices  $C_p^\perp(y)$  of the cap projection  $C^\perp$  by sweeping a line parallel to the  $\bar{x}$ -axis. A point  $(x, y)$  inside  $C^\perp$  is sampled by first finding the  $y$  value that cuts a slice with a desired partial area  $A(C_p^\perp(y))$  and then uniformly sampling a coordinate  $x$  between the intersection coordinates  $x_{\min}$  and  $x_{\max}$ .

#### 4.1. Parallel map

Our parallel map considers slices of  $C^\perp$  enclosed between the  $\bar{x}$ -axis and a chord parallel to that axis, as illustrated in Figure 3. To sample a point inside  $C^\perp$ , we first find the chord that cuts a slice with a desired area and then choose a point along that chord with uniform density.

A slice  $C_p^\perp(y)$  is cut by a chord with endpoints  $(x_{\min}, y)$  and  $(x_{\max}, y)$ , with  $x_{\min} < x_{\max}$  and  $y \in [0, a_y]$ . The area of such a slice is the sum of the areas of differential rectangles of width  $(x_{\max} - x_{\min})$  and height  $dy$ :

$$A_p(y) = A(C_p^\perp(y)) = \int_0^y [x_{\max}(y') - x_{\min}(y')] dy'. \quad (12)$$

For a given  $y$ , to compute the slice area  $A_p(y)$  we need expressions for the chord endpoint coordinates  $x_{\min}$  and  $x_{\max}$ . To find these, we use the ellipse and circle equations as in Equation (10). The point  $(x_{\min}, y)$  is always on the ellipse, so we have  $E(x_e - x_{\min}, y) = 1$ . The point  $(x_{\max}, y)$  is on the ellipse when  $\alpha \leq \beta$  or when  $y_l < y$ , with  $E(x_{\max} - x_e, y) = 1$ , and on the unit circle when  $|\beta| < \alpha$  and  $y \leq y_l$ , with  $C(x_{\max}, y) = 1$ . From these we obtain the sought expressions:

$$x_{\min}(y) = x_e - a_x \sqrt{1 - \frac{y^2}{a_y^2}}, \quad (13)$$

$$x_{\max}(y) = \begin{cases} x_e + a_x \sqrt{1 - \frac{y^2}{a_y^2}}, & \text{if } \alpha \leq \beta \text{ or } y_l < y, \\ \sqrt{1 - y^2}, & \text{if } |\beta| < \alpha, y \leq y_l. \end{cases} \quad (14)$$

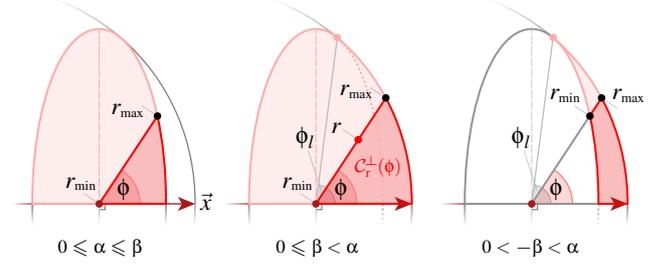
Integrating these expressions, for  $A_p$  we get:

$$A_p(y) = \begin{cases} 2A_{pE}(y), & \text{if } 0 \leq \alpha \leq \beta, \\ 2A_{pE}(y) + A_{pC}(y'') - A_{pE}(y''), & \text{if } 0 \leq \beta < \alpha, \\ A_{pE}(y'') + A_{pC}(y''), & \text{if } 0 < -\beta < \alpha, \end{cases} \quad (15)$$

where  $y'' = \min(y, y_l)$  and

$$A_{pE}(v) = a_x a_y I\left(\frac{v}{a_y}, 1\right), \quad A_{pC}(v) = I(v, 1) - x_e v. \quad (16)$$

Note that in the lune-only case, since  $A_{pE}$  is negative, it holds that



**Figure 4:** Our radial map cuts sectors  $C_r^\perp(\phi)$  of the cap projection  $C^\perp$  by sweeping a radial line about the ellipse center. A point  $(r, \phi)$  inside  $C^\perp$  is sampled by first finding the angle  $\phi$  that cuts a sector with a desired partial area  $A(C_r^\perp(\phi))$  and then sampling a radius  $r$  with linear density between the intersections  $r_{\min}$  and  $r_{\max}$ .

$A_p(y) = A_p(y_l)$  whenever  $y > y_l$ . In that case the range of  $y$  is  $[0, y_l]$ , as seen in Figure 3, left. In Equation (16) above we use

$$I(u, w) = \frac{1}{2} \left[ w u \sqrt{1 - u^2} + \arcsin(u) \right]. \quad (17)$$

**Map evaluation.** Having characterized the slice area, we can map a given unit-square point  $(s, t)$  to a point  $(x, y)$  in  $C^\perp$ . We first need to find the coordinate  $y$  that satisfies  $A_p(y) = t A_p(a_y)$ , for which we need to evaluate the inverse function  $A_p^{-1}$ . An analytical form of this inverse is not known, therefore we resort to numerical root finding (in the positive half of  $C^\perp$ ), which yields a non-negative value  $y^*$ . Then we use the sampling region symmetry to obtain  $y$  as

$$y = \begin{cases} -y^* \Leftarrow A_p(y^*) - (1 - 2t)A_p(a_y) = 0, & \text{if } t < 0.5, \\ y^* \Leftarrow A_p(y^*) - (2t - 1)A_p(a_y) = 0, & \text{if } t \geq 0.5, \end{cases} \quad (18)$$

which covers both the negative and positive halves of  $C^\perp$ . Once we have  $y$ , the chord endpoint coordinates  $x_{\min}$  and  $x_{\max}$  are given by Equations (13) and (14). The  $x$  coordinate of the sought point is finally obtained via linear interpolation between these endpoints:

$$x = x_{\min} + s(x_{\max} - x_{\min}). \quad (19)$$

Finally, by plugging  $x$  and  $y$  into Equation (7) we obtain the corresponding world-space direction on the spherical cap.

#### 4.2. Radial map

Our radial map considers sectors of  $C^\perp$  enclosed between the  $\bar{x}$ -axis and a radial half-line rotating about the ellipse center, as illustrated in Figure 4. To sample a point inside  $C^\perp$ , we first find the radial line that cuts a sector with a desired area and then choose a radius along that line with linear density.

The half-line makes angles  $\phi \in [0, \pi]$  with the  $\bar{x}$ -axis. For some radius  $r$ , a differential angle  $d\phi$  around the line forms a differential triangle with area  $r^2 d\phi/2$ . The differential region between the radii  $r_{\min}$  and  $r_{\max}$  of intersection between the line and  $C^\perp$  is then a truncated triangle, i.e. a trapezoid, with area  $(r_{\max}^2 - r_{\min}^2) d\phi/2$ . We are interested in the sector  $C_r^\perp(\phi)$  that the half-line cuts at angle  $\phi$ , whose area is a sum of such differential areas:

$$A_r(\phi) = A(C_r^\perp(\phi)) = \frac{1}{2} \int_0^\phi [r_{\max}^2(\phi') - r_{\min}^2(\phi')] d\phi'. \quad (20)$$

To compute the sector area  $A_r(\phi)$ , we need expressions for the intersection radii  $r_{\min}$  and  $r_{\max}$ , for which we again use the ellipse and circle equations as in Equation (10). The point at  $r_{\min}$  is at the ellipse center if  $0 \leq \beta$ , and on the ellipse otherwise, with  $E(r_{\min} \cos \phi, r_{\min} \sin \phi) = 1$ . The point at  $r_{\max}$  is on the ellipse when  $\alpha \leq \beta$  or when  $\phi_l < \phi$ , with  $E(r_{\max} \cos \phi, r_{\max} \sin \phi) = 1$ , and on the unit circle when  $|\beta| < \alpha$  and  $\phi \leq \phi_l$ , where  $\phi_l = \arctan(y_l/(x_l - x_e))$ , with  $C(x_e + r_{\max} \cos \phi, r_{\max} \sin \phi) = 1$ . From these we obtain the sought expressions:

$$r_{\min}(\phi) = \begin{cases} 0, & \text{if } 0 \leq \beta, \\ \frac{a_x}{\sqrt{1 - \cos^2 \beta \sin^2 \phi}}, & \text{if } 0 < -\beta, \end{cases} \quad (21)$$

$$r_{\max}(\phi) = \begin{cases} \frac{a_x}{\sqrt{1 - \cos^2 \beta \sin^2 \phi}}, & \text{if } \alpha \leq \beta \text{ or } \phi_l \leq \phi, \\ \sqrt{1 - x_e^2 \sin^2 \phi} - x_e \cos \phi, & \text{if } |\beta| < \alpha, \phi < \phi_l. \end{cases} \quad (22)$$

Integrating these expressions, for  $A_r$  we get:

$$A_r(\phi) = \begin{cases} A_{rE}(\phi), & \text{if } 0 \leq \alpha \leq \beta, \\ A_{rE}(\phi) + A_{rC}(\phi'') - A_{rE}(\phi''), & \text{if } 0 \leq \beta < \alpha, \\ A_{rE}(\phi'') + A_{rC}(\phi''), & \text{if } 0 < -\beta < \alpha, \end{cases} \quad (23)$$

where  $\phi'' = \min(\phi, \phi_l)$ . Similarly to the parallel map, we effectively clamp  $\phi \in [0, \phi_l]$  in the lune-only case, as seen necessary in Figure 4, right. Above we use

$$A_{rE}(\phi) = \frac{1}{2} a_x a_y \arctan\left(\frac{a_x}{a_y} \tan \phi\right), \quad (24)$$

$$A_{rC}(\phi) = I(\sin \phi, x_e^2) - I(x_e \sin \phi, 1), \quad (25)$$

where  $I(u, w)$  is defined in Equation (17) above.

**Map evaluation.** Having characterized the sector area, we can map a given unit-square point  $(s, t)$  to a point with polar coordinates  $(r, \phi)$  in  $\mathcal{C}^\perp$ . We first need to find the angle  $\phi$  that satisfies  $A_r(\phi) = t A_r(\pi)$ . Since an analytical form of the inverse function  $A_r^{-1}$  is not known in the general case, we resort to numerical root finding (in the positive half of  $\mathcal{C}^\perp$ ), which yields a non-negative angle  $\phi^*$ . We again use the sampling region symmetry to obtain  $\phi$ :

$$\phi = \begin{cases} -\phi^* \Leftarrow A_r(\phi^*) - (1 - 2t) A_r(\pi) = 0, & \text{if } t < 0.5, \\ \phi^* \Leftarrow A_r(\phi^*) - (2t - 1) A_r(\pi) = 0, & \text{if } t \geq 0.5, \end{cases} \quad (26)$$

which covers both the negative and positive halves of  $\mathcal{C}^\perp$ . Once we have  $\phi$ , the radii  $r_{\min}$  and  $r_{\max}$  are given by Equations (21) and (22). The radius  $r$  of the sought point is then located between  $r_{\min}$  and  $r_{\max}$  with linear density, which accounts for the triangular shape of the differential sampling region:

$$r = \sqrt{r_{\min}^2 + s(r_{\max}^2 - r_{\min}^2)}. \quad (27)$$

As seen in Figure 4, unless  $\beta < 0$ , we have  $r_{\min} = 0$  in which case sampling simplifies to  $r = r_{\max} \sqrt{s}$  as in traditional polar disk mapping [Dut03, p. 13]. The Cartesian coordinates of the sampled point in the local frame  $\mathcal{R}$  are then (see Figure 2)

$$x = x_e + r \cos \phi, \quad y = r \sin \phi. \quad (28)$$

Finally, by plugging  $x$  and  $y$  into Equation (7) we obtain the corresponding world-space direction on the spherical cap.

**Ellipse-only case.** A useful sampling optimization opportunity arises in the ellipse-only case, i.e. when  $\alpha \leq \beta$ . This is in fact the most commonly occurring case in practical rendering scenarios, where the projected cap area simplifies to  $A_r(\phi) = A_{rE}(\phi)$ . The expression for  $A_{rE}(\phi)$  has a geometrical interpretation: it is the area of an elliptical sector, which in turn is a scaled unit-radius disk sector. The scaling factors are the ellipse's semi-minor and semi-major axes,  $a_x$  and  $a_y$ . Thus, in the ellipse-only case we can use the simple analytical polar disk mapping to first sample  $(r, \phi) = (\sqrt{s}, 2\pi t)$ , and the corresponding Cartesian coordinates can then be computed as  $(x, y) = (x_e + a_x r \cos \phi, a_y r \sin \phi)$ . This optimization makes the sampling faster and more accurate numerically, and also preserves the continuity of the map w.r.t. continuous changes in  $\alpha$  and  $\beta$ .

### 4.3. Numerical inversion

Evaluating our maps requires computing a distance  $y = A_p^{-1}(u)$  or an angle  $\phi = A_r^{-1}(u)$  for some given partial area  $u < A(\mathcal{C}^\perp)/2$ . Since there are no analytical expressions for either of these inverse functions in the general case, we must resort to numerical inversion. We use the Newton-Raphson iterative root finding method as it has a quadratic convergence rate. In some cases, when the derivative of the inverted function is small, the Newton method can cause the estimated solution to go out of range; in that case we track the current known solution interval and do a binary split of that interval.

### 4.4. Differential form factor

As a side result of our derivations, we obtain closed-form expressions for the form factor  $F$  between a differential surface element and a sphere. The form factor is the ratio of the area of the projected spherical cap  $\mathcal{C}^\perp$  and the area  $\pi$  of the unit-radius disk. This can be computed in two ways, using either  $A_p$  from Equation (15) or  $A_r$  from Equation (23):

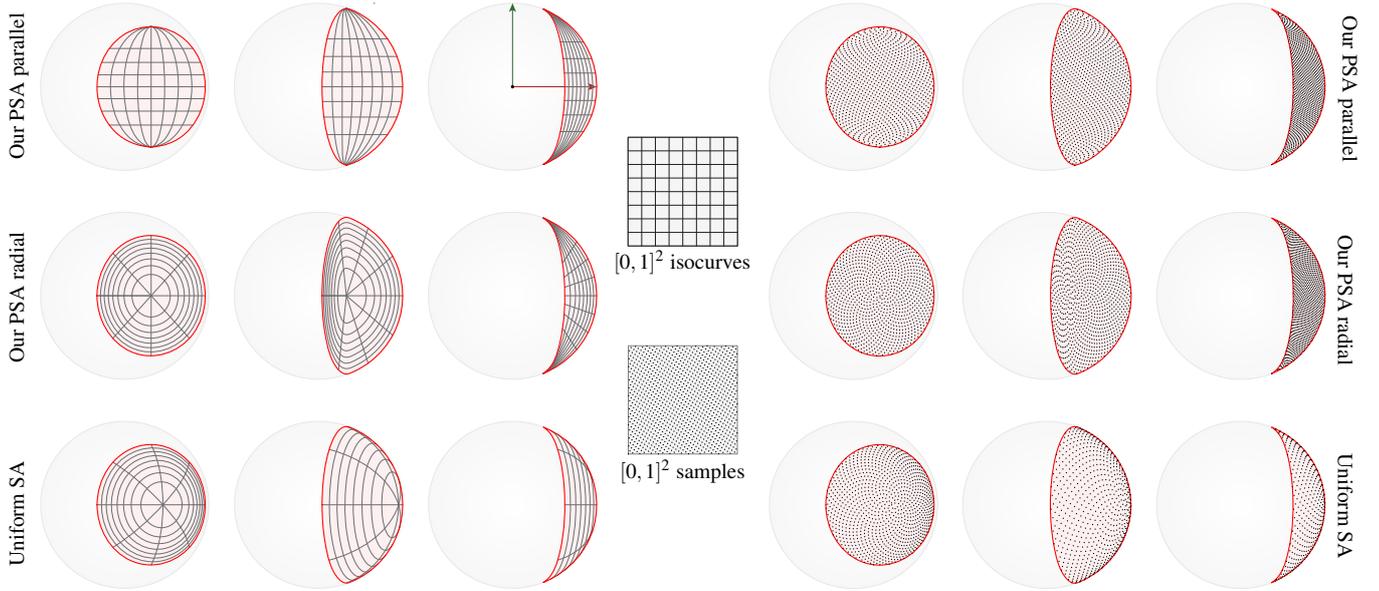
$$F = \frac{A(\mathcal{C}^\perp)}{\pi} = \frac{2A_p(a_y)}{\pi} = \frac{2A_r(\pi)}{\pi}. \quad (29)$$

While the form factor expressions for fully visible spherical polygons [BRW89] and spherical caps [Hei17] are simpler, the above expressions also cover the more general case where the cap can also be partially below the horizon. Our expressions in Equation (29) are equivalent to that of Snyder [Sny96].

The cap projection area  $A(\mathcal{C}^\perp)$  is also needed when estimating the scattered radiance integral  $L_o$  using our maps. In Equation (5) it appears explicitly, whereas in Equation (2) it appears in the pdf of the sampled direction, as we discussed at the end of Section 2.1.

## 5. Implementation and results

In this section we discuss the implementation and evaluate the qualities of our two projected solid angle (PSA) preserving maps against the traditional solid angle (SA) preserving map used for uniform spherical cap sampling [Wan92]. We compare the warping properties of the three maps both visually and for rendering the illumination from spherical light sources.



**Figure 5:** A visual comparison between our projected solid angle (PSA) maps and the standard solid angle (SA) map. We plot canonical  $[0, 1]^2$  isocurves (left) and a 1024-point rank-1 lattice (right) warped using each map, and show only those falling in the upper hemisphere. Note the non-uniformity in the projected distribution produced by the SA map, which also places most samples below the horizon in the lune-only case. Our PSA maps always generate samples in the visible part of the spherical cap, i.e. only directions in the upper hemisphere.

### 5.1. Implementation

Listing 1 contains code for our radial map, which we have found to be faster to evaluate than our parallel map thanks to its efficient handling of the (most common) ellipse-only case. The code operates in the local frame  $\mathcal{R}$  described in Section 3.1. We exclude optimizations for the sake of keeping the code compact. The full supplemental source code includes such optimizations [UnG18].

### 5.2. Map evaluation

In Figure 5, we show canonical  $[0, 1]^2$  isocurves and a sample point set transformed through the three maps. Isocurve portions and samples below the horizon are not plotted. The point set is a rank-1 lattice with 1024 points created using the Korobov rule [Nie92, p. 124] with generating vector  $(1, 275)$ . We chose a lattice as it more clearly shows distortions introduced by the maps.

The isocurve plots clearly show that the SA map yields cells of different size, while our PSA maps preserve relative areas by producing equally-sized cells, as expected. The projected density of the warped point set is thus uniform with our maps, unlike with the SA map. Moreover, when the spherical cap is partially below the horizon (the ellipse+lune and lune-only cases) the SA map generates invalid samples in the lower hemisphere. In Figure 5 this is reflected as having fewer points in the SA map plots in those cases. In contrast, our PSA maps are guaranteed to place all the samples in the upper hemisphere, i.e. in the visible part of the spherical cap.

Since PSA maps cannot be conformal, i.e. angle-preserving, cell distortions cannot be avoided. Nevertheless, the discrepancy and inter-point distance properties of the canonical input lattice are well preserved by our maps in the example shown in Figure 5.

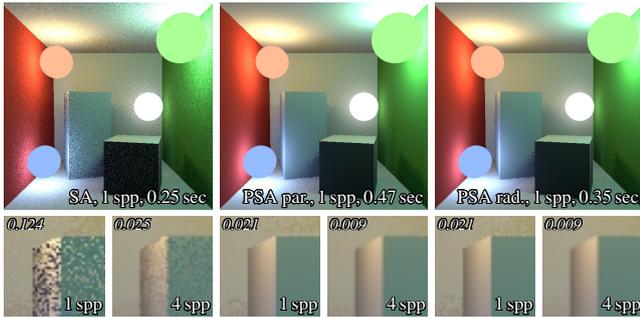
```
// Evaluates the integrand of  $A_r$  (20);  $\phi$  must be in  $[0, \pi]$ 
real RadIntegrand( real phi ) // real can be float or double
{
    real rmi= rmin(phi), rma= rmax(phi); //  $r_{\min}$  (21) and  $r_{\max}$  (22)
    return 0.5*( rma*rma - rmi*rmi ); // see (20)
}

// Evaluates  $\phi = A_r^{-1}(\text{area})$ ;  $\text{area}$  must be in  $[0, A_r(\pi)]$ , returned angle  $\phi$  is in  $[0, \pi]$ 
real ArInverse( real area )
{
    RealFunc F = Ar, // function being inverted ( $A_r$ )
    f = RadIntegrand; // derivative of  $F$  (integrand)
    const real arg_max = M_PI, // max. argument of  $F$  (min. is 0)
    target = area, // argument of  $F^{-1}$ 
    target_max = F( arg_max ); // max. value of  $F$  (min. is 0)
    // InverseNSB performs numerical inversion, not shown here
    return InverseNSB( F, f, arg_max, target, target_max );
}

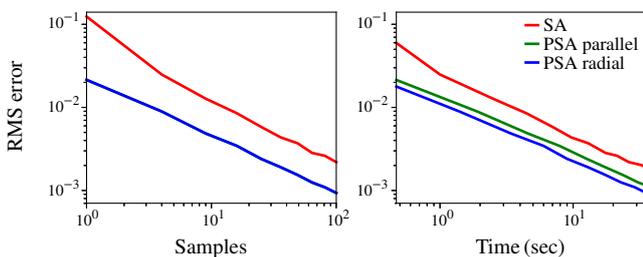
// Evaluates the radial map: computes  $(\mathbf{x}, \mathbf{y}) \in \mathcal{C}^\perp$  as a function of  $(\mathbf{s}, \mathbf{t}) \in [0, 1]$ 
void RadialMap( real s, real t, real &x, real &y )
{
    const bool fv = alpha <= beta; // true iff spherical cap is fully visible
    const real
        u = t < 0.5 ? 1.0-2.0*t : 2.0*t -1.0,
        phi = fv ? M_PI*u : ArInverse( u*0.5*Ar( M_PI ) ),
        rmi = fv ? 0.0 : rmin( phi ), //  $r_{\min}$  (21)
        rma = fv ? 1.0 : rmax( phi ), //  $r_{\max}$  (22)
        si = t < 0.5 ? -sin( phi ) : sin( phi ),
        co = sqrt( 1.0-si*si )*( phi <= M_PI*0.5 ? 1.0 : -1.0 ),
        rad = sqrt( s*rma*rma + (1.0-s)*rmi*rmi );
    x = fv ? xe + ax*rad*co : xe + rad*co;
    y = fv ? ay*rad*si : rad*si;
}

```

**Listing 1:** Evaluation of our radial map. Note that the code avoids numerical inversion when the spherical cap is fully visible.



**Figure 6:** Scene rendered with the standard SA map (left) and our PSA parallel (middle) and PSA radial (right) maps with 1 and 4 samples/pixel (spp). In the top-left corner of each zoom-in we report the RMS error w.r.t. a 1024-sample SA-map reference image.



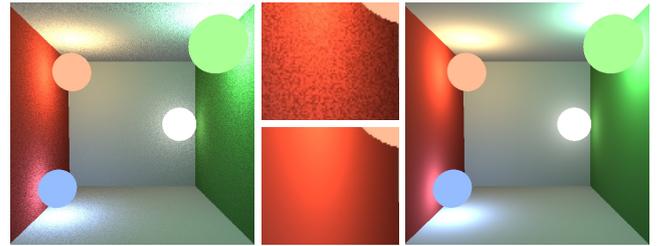
**Figure 7:** RMS error plots of the three images on the top row of Figure 6 as functions of the number of samples (left) and the render time (right). Our two PSA maps yield identical error for the same number of samples, thus their RMS plots on the left overlap.

### 5.3. Rendering evaluation

We have implemented our two PSA maps in the Arnold production renderer as drop-in replacements for the traditional uniform SA sampling of spherical light sources. In the rendering benchmarks presented below we always take one sample per light source per camera ray. Taking multiple samples per light source would play in favor of our PSA techniques, since they have a higher sampling initialization cost than uniform SA sampling.

Figure 6 shows a purely diffuse Cornell Box scene with four light sources that do not act as light occluders, rendered using all three techniques with 1 and 4 samples per pixel. Our PSA techniques bring most variance reduction on surface regions illuminated from grazing angles or with large subtended solid angle where there is a large variation in the cosine foreshortening term. On this simple scene our parallel and radial maps increase the render times by about 90% and 40%, respectively. This is due to the more involved sampling code which also includes numerical inversion. The radial map is the faster of the two thanks to the efficient disk-based sampling in the ellipse-only case (see Section 4.2 and Listing 1). The increased render times are more than compensated for by the reduction in the root mean squared (RMS) error.

In Figure 7 we show log-log RMS error plots of the three images in Figure 6 as functions of the number of samples and the render time. As expected, our maps yield lower error for the same num-



**Figure 8:** In the absence of occlusion, solid angle sampling (left) has non-zero variance, while our method (right) gives the exact illumination solution with a single sample per light source.

ber of samples (identical for the two maps, hence the overlapping plots). They are also more computationally efficient as they achieve lower error for the same render time. The radial map is consistently faster on this scene where the light sources are fully visible from most shading points.

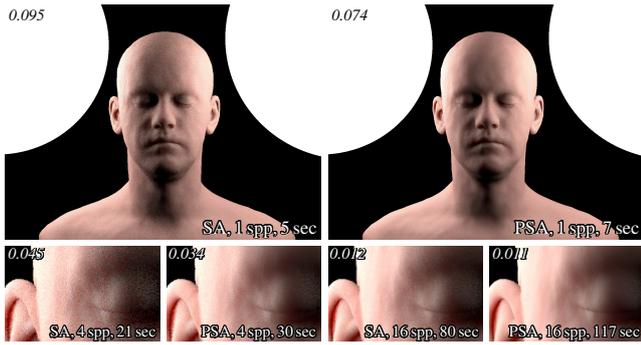
Removing all light blockers from the Cornell box scene, in Figure 8 we confirm that in the absence of occlusion and with uniform light emission, our PSA method yields the exact, noise-free direct illumination solution, even with a single sample per pixel. All noise in the left image is a result of the imperfect direction distribution produced by the uniform SA sampling.

The head model shown in Figure 9 has a material that simulates subsurface scattering by sampling nearby surface points at every shading point and averaging their direct illumination weighted by a diffusion profile. While this additional sampling step prevents our PSA method from achieving zero-variance estimation, it can sample the grazing-angle illumination from the two spherical lights more efficiently than uniform SA sampling.

The scene shown in Figure 10 is illuminated by a setting sun-like distant light source that is mostly below the horizon. The images are dominated by shadow noise which our PSA sampling does not address. Nevertheless our method still brings noticeable improvement on surfaces with normal almost perpendicular to the light source direction, such as the ground, for a modest increase in render time.

## 6. Conclusion

In this paper we present practical techniques for stratified sampling of the projected solid angle (PSA) subtended by a spherical luminaire at a surface point. Our techniques allow to completely remove noise for diffuse reflectors, in the absence of occluders, even if the luminaire is partially below the surface tangent plane. In the presence of occluders, our techniques yield lower noise compared to standard solid angle (SA) sampling which does not account for the cosine term in the radiance contribution. Although our PSA sampling is more costly than SA sampling, it is more computationally efficient and its overhead becomes negligible on complex scenes where sample evaluation is much more costly than sample generation. PSA sampling reduces variance in general, and in practice we observe this reduction in our scenes, although in some particular cases, when the cosine term is negatively correlated with the visibility term, PSA sampling may increase variance. This is a po-



**Figure 9:** Subject illuminated by two spherical lights, rendered with direct illumination and diffusion-profile subsurface scattering. The variance reduction brought by our PSA sampling (radial map) over uniform SA sampling is substantial in regions where the light source is mostly below the surface tangent plane.

tential pitfall of any importance sampling technique whose distribution does not account for the entire integrand.

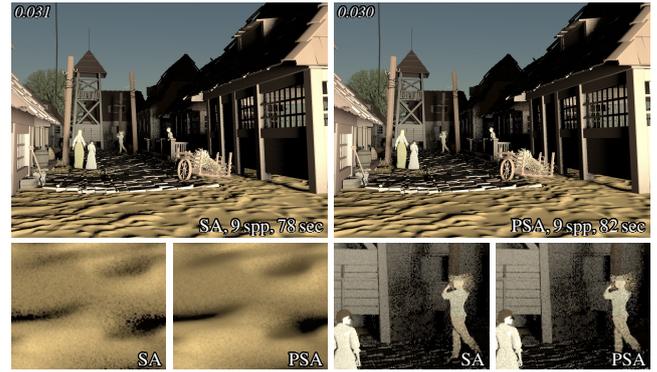
Our method can be used to produce a sample distribution with clamped-cosine density on an arbitrary spherical cap. Using the method of Dupuy et al. [DHB17], this source distribution can in turn be transformed onto a given destination spherical cap, in order to fit a target distribution on that cap. In addition, our formulations provide novel closed-form expressions for the form factor between a differential surface element and a sphere [Sny96]. The form factor could be used to improve many-light rendering techniques [HKWB09, SHD15, EK17] as well as direct illumination sampling with combined analytical and stochastic techniques [HHM18].

### Acknowledgments

The authors thank Marcos Fajardo and the anonymous reviewers for their constructive feedback. Carlos Ureña acknowledges funding from Autodesk for this work.

### References

- [Arv95] ARVO J.: Stratified sampling of spherical triangles. In *Proceedings of SIGGRAPH '95* (1995), ACM, pp. 437–438. 1
- [Arv01] ARVO J.: Stratified sampling of 2-manifolds. *SIGGRAPH 2001 Course Notes* 29, 2 (2001). 1
- [BRW89] BAUM D. R., RUSHMEIER H. E., WINGET J. M.: Improving radiosity solutions through the use of analytically determined form-factors. *SIGGRAPH Comput. Graph.* 23, 3 (July 1989), 325–334. 5
- [CWH93] COHEN M. F., WALLACE J., HANRAHAN P.: *Radiosity and Realistic Image Synthesis*. Academic Press Professional, Inc., San Diego, CA, USA, 1993. 1
- [DHB17] DUPUY J., HEITZ E., BELCOUR L.: A spherical cap preserving parameterization for spherical distributions. *ACM Trans. Graph.* 36, 4 (July 2017), 139:1–139:12. 1, 8
- [Dut03] DUTRE P.: *Global illumination compendium*, 2003. 5
- [EK17] ESTEVEZ A. C., KULLA C.: Importance sampling of many lights with adaptive tree splitting. In *ACM SIGGRAPH 2017 Talks* (New York, NY, USA, 2017), SIGGRAPH '17, ACM, pp. 33:1–33:2. 8



**Figure 10:** Scene illuminated by a sun-like infinite light source low on the horizon. Our PSA sampling (parallel map) brings improvement over SA sampling in regions illuminated from grazing angles, but cannot help in regions dominated by shadow noise.

- [Gam16] GAMITO M. N.: Solid angle sampling of disk and cylinder lights. *Computer Graphics Forum* 34, 4 (2016). 1
- [GUnK\*17] GUILLÉN I., UREÑA C., KING A., FAJARDO M., GEORGIEV I., LÓPEZ-MORENO J., JARABO A.: Area-preserving parameterizations for spherical ellipses. *Computer Graphics Forum* 36, 4 (2017), 179–187. 1
- [Hei17] HEITZ E.: *Geometric Derivation of the Irradiance of Polygonal Lights*. Research report, Unity Technologies, Jan. 2017. 5
- [HHM18] HEITZ E., HILL S., MCGUIRE M.: Combining analytic direct illumination and stochastic shadows. In *ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games, Montreal, Quebec, Canada: 15-18 May 2018 (short paper)* (2018). 8
- [HKWB09] HAŠAN M., KRIVÁNEK J., WALTER B., BALA K.: Virtual spherical lights for many-light rendering of glossy scenes. *ACM Trans. Graph.* 28, 5 (dec 2009), 143:1–143:6. 8
- [Nie92] NIEDERREITER H.: *Random Number Generation and quasi-Monte Carlo Methods*. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA, 1992. 6
- [PJH16] PHARR M., JAKOB W., HUMPHREYS G.: *Physically Based Rendering: From Theory to Implementation*, 3rd ed. Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2016. 1
- [SHD15] SIMON F., HANIKA J., DACHSBACHER C.: Rich-VPLs for improving the versatility of many-light methods. *Computer Graphics Forum* 34, 2 (2015), 575–584. 8
- [Sny96] SNYDER J. M.: *Area Light Sources for Real-Time Graphics*. Research report, Microsoft Research, March 1996. 5, 8
- [SWZ96] SHIRLEY P., WANG C., ZIMMERMAN K.: Monte Carlo techniques for direct lighting calculations. *ACM Trans. Graph.* 15, 1 (jan 1996), 1–36. 1, 2
- [UFK13] UREÑA C., FAJARDO M., KING A.: An area-preserving parameterization for spherical rectangles. *Computer Graphics Forum* 32, 4 (2013), 59–66. 1
- [Un00] UREÑA C.: Computation of irradiance from triangles by adaptive sampling. *Computer Graphics Forum* 19, 2 (2000), 165–171. 1
- [UnG18] UREÑA C., GEORGIEV I.: Projected spherical cap sampling open source code repository. <https://github.com/carlos-urena/psc-sampler>, 2018. 6
- [Wan92] WANG C.: Physically correct direct lighting for distribution ray tracing. In *Graphics Gems III* (1992), Academic Press Professional, Inc., pp. 307–313. 1, 2, 5