

Real-time Approximation of Photometric Polygonal Lights

CHRISTIAN LUKSCH, VRVis Research Center

LUKAS PROST, rmDATA

MICHAEL WIMMER, TU Wien



Fig. 1. Photometric point lights (left), a typical simplification of photometric area lights used in real-time rendering systems. Our real-time technique (middle) provides a well-defined illumination in the near-field of the light sources, closely resembling the reference solution (right) rendered with 5000 point lights.

We present a real-time rendering technique for photometric polygonal lights. Our method uses a numerical integration technique based on a triangulation to calculate noise-free diffuse shading. We include a dynamic point in the triangulation that provides a continuous near-field illumination resembling the shape of the light emitter and its characteristics. We evaluate the accuracy of our approach with a diverse selection of photometric measurement data sets in a comprehensive benchmark framework. Furthermore, we provide an extension for specular reflection on surfaces with arbitrary roughness that facilitates the use of existing real-time shading techniques. Our technique is easy to integrate into real-time rendering systems and extends the range of possible applications with photometric area lights.

CCS Concepts: • **Computing methodologies** → **Rendering**.

Additional Key Words and Phrases: real-time rendering, area lights, photometric lights

ACM Reference Format:

Christian Luksch, Lukas Prost, and Michael Wimmer. 2020. Real-time Approximation of Photometric Polygonal Lights. *Proc. ACM Comput. Graph. Interact. Tech.* 3, 1, Article 4 (May 2020), 18 pages. <https://doi.org/10.1145/3384537>

1 INTRODUCTION

Accurate simulation of illumination from real-world luminaires is an essential requirement in lighting design and architectural planning. In simulation tools, these luminaires are typically referred to as *Photometric Lights*. Their emission characteristics are provided by photometric measurement data and a geometric primitive, such as a polygon or disk, defining the region of emission. Typically, offline rendering algorithms are used in this context, including indirect

Authors' addresses: Christian Luksch, VRVis Research Center; Lukas Prost, rmDATA; Michael Wimmer, TU Wien.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than the author(s) must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2020 Copyright held by the owner/author(s). Publication rights licensed to ACM.

2577-6193/2020/5-ART4 \$15.00
<https://doi.org/10.1145/3384537>

illumination with complex light interactions. However, this makes work processes slow and limits interactivity. Real-time rendering systems with physically based shading are starting to overcome this restriction. They provide very convincing image quality with support of photometric point lights [Donzallaz 2019] that are already widely used for architectural visualization. While there are numerous real-time solutions for surfaces with diffuse emission [Drobot 2014; Heitz et al. 2016; Lecocq et al. 2016], area lights with direction-dependent intensity profiles remain a challenge for real-time rendering.

In this paper, we present a real-time approximation of photometric polygonal lights providing a well-defined near-field illumination depicting the shape of the light emitter and its characteristics (Figure 1). Our technique (middle) uses a numerical integration technique based on a triangulation to calculate the diffuse shading at reasonable additional costs compared to the point representation (left). We introduce an optimized sampling strategy, combining a dynamic point with fixed sample positions to build a triangulation of the integration domain. This provides suitable sample weights for the integration and a continuous illumination while avoiding sampling artifacts. We evaluate the accuracy and performance using a diverse selection of photometric data sets. Furthermore, we provide an extension for specular reflection on microfacet surfaces by deriving an energy-equivalent diffuse emitter as substitution and applying the *Linearly Transformed Cosines* (LTC) [Heitz et al. 2016] method. Our technique can be easily integrated into existing real-time rendering systems and enables convincing noise-free shading with polygonal photometric lights. Previously this was only possible using offline rendering methods.

2 BACKGROUND

In this section, we give an introduction to photometric lights and their representation in a rendering system. Next, we summarize related work and recent advances in the field of rendering algorithms and real-time approximations of area lights.

2.1 Photometric Lights

In lighting design and illumination engineering, photometry is used to evaluate, describe and characterize the photometric performance of luminaires. Far-field photometry is the current industry standard and virtually all commercially available data is acquired using this procedure [DiLaura et al. 2011, 9.24]. Measurements of luminous intensities are taken from relatively large distances and described as a spherical function $I(\theta, \phi)$. This data can be used to calculate the illumination, created by a luminaire, as anisotropic point light source. For diffuse emitters it is proven that if the distance is greater than five times the maximum dimension of the emitter the error is less than 2%. This “five times rule” has been adopted for practical applications of far-field photometry [DiLaura et al. 2011, 10.10]. In rendering and lighting design software far-field photometric data can be imported using IES or LDT files. In addition to the luminous intensity distribution, these files provide information on the size of the luminaire and the size of the luminous area. Including the geometric shape allows illustration of a plausible near-field illumination. Verbeck and Greenberg [1984] presented an early system following this workflow, where light-emitting surfaces are evaluated as a collection of anisotropic point lights. A more general system and elaborated rendering kernel in this context is Radiance [Ward 1994]. In order to allow validation of lighting design tools using photometric data, the Commission International de l’Eclairage (CIE) has provided a collection of test cases [Ashdown et al. 2006].

Accurate calculation for short distances requires near-field photometry. In particular, there are two classes of methods. Application-distance photometry models the luminaire as a set of measurements from several distances and uses the same procedures as with far-field photometry. However, this kind of data cannot be extrapolated to planes that are tilted with respect to those measured. In

order to overcome this limitation, Ashdown [1993; 1995] introduced luminance-field photometry, where the field of light surrounding a luminaire is measured. Using this data, the illuminance can be calculated at any point, but the size of the data can be difficult to manage [Ashdown and Rykowski 1998]. This type of measurement data is typically available for light sources (lamps, LEDs, etc.) and used by luminaire manufacturers to design optics and reflectors. However, near-field photometry is not used in practice in the field of lighting design due to its complexity and is therefore not part of our research agenda.

2.2 Related Work

Monte Carlo methods. Distributed ray tracing [Cook et al. 1984] and path tracing [Kajiya 1986] are the most general algorithms used in computer graphics and allow the calculation of global illumination and fuzzy phenomena, such as soft shadows, motion blur and depth of field. In these algorithms area lights with arbitrary intensity profiles or textured surfaces can be described in an inherent way. In general, light transport can be formulated either as an integral over solid angles or over all surfaces [Shirley et al. 1996]. Recently, several reviewed and new specialized methods for sampling of rectangles [Ureña et al. 2013], disks and ellipses [Guillén et al. 2017] and spherical caps [Peters and Dachsbacher 2019; Ureña and Georgiev 2018] have been presented. Techniques to distribute samples according to the product of two functions, such as used in image-based lighting [Burke et al. 2005; Clarberg et al. 2005; Estevez and Lecocq 2018], could also be applied to photometric lights.

Using Monte Carlo rendering algorithms in real-time is very challenging as only very low sample budgets can be used per frame resulting in heavy noise. Nevertheless, advances in GPU accelerated ray tracing and new de-noising techniques exploiting temporal re-projection [Mara et al. 2017; Schied et al. 2017] or extended filter domains, such as texture [Munkberg et al. 2016] or path space filtering [Binder et al. 2019], extend the horizon of feasible real-time applications. Hybrid solutions can also provide otherwise hard to achieve effects at reasonable costs. Heitz et al. [2018] combined analytic area light illumination with sampled shadows and list other possible applications.

An alternative rendering algorithm that is similarly well suited to render arbitrary area lights is *Instant Radiosity* [Keller 1997]. It is the foundation of many-light rendering methods and there is a wide variety of improvements. We refer to the survey of Dachsbacher et al. [2014] for a detailed summary. Advantages are that this algorithm is not susceptible to noise and is very scalable, making it interesting for real-time applications [Hedman et al. 2016; Laurent et al. 2016; Lin and Yuksel 2019; Luksch et al. 2019].

Point approximations. Assuming full visibility, diffuse emission and diffuse receiver an area light integral can be sampled using a single point. This point is called the *Most Representative Point* (MRP). Although the optimal point can only be found using an expensive calculation, simplifications provide plausible results. Based on this concept, Wang et al. [2008] developed their approach for textured lights. Recent techniques [Drobot 2014; Karis 2013], integrated into real-time rendering systems with physically based shading, provide meaningful normalization and include specular reflections. Lagarde and Rousiers [2014] propose *Structured Sampling* as a more accurate approximation. Adopting such a technique for photometric lights might be able to overcome the limitations of fixed-point representations. In Section 4 we will share our experiments on adopting these approximations to photometric lights.

Analytic solutions. Certain geometric configurations allow light transport calculation with exact analytic solutions. A lot of methods have been adopted from the field of radiative heat transfer analysis and applied to *Radiosity* [Goral et al. 1984] calculations. There the diffuse radiative energy transfer between surfaces is described by so-called form factors. Baum et al. [1989] use an analytic

solution for the transfer from polygon to point. This method has found applications in numerous real-time techniques [Heitz et al. 2016; Luksch et al. 2013; McGuire 2010]. Lagarde and Rousiers [2014] summarize other practical analytic solutions for disks and spheres.

While the above solutions are limited to diffuse surfaces, Arvo et al. [1995] introduced Irradiance Tensors that allows formulating non-diffuse transport functions. However, their solution is restricted to the Phong distribution and has impractical runtime complexity. Lecocq et al. [2016] solve these drawbacks and provide a real-time shading technique for surfaces with microfacet BRDFs lit by polygonal area lights. Heitz et al. [2016; 2017] use a different approach and apply a pre-processing step to approximate area light shading using *Linearly Transformed Cosines* (LTC) that can be evaluated analytically in real-time. As a complete diffuse and specular light transport is important for consistent renderings, we will provide an extension to our photometric light approximation based on this method for specular reflection.

Complex light sources. Heidrich et al. [1998] introduced *Canned Light Sources* to speed up global illumination with complex light sources by substituting them with a pre-computed light field. Such an approach is also related to the application of near-field photometry. Mas et al. [2008] presented a compression and importance sampling method for near-field measurements registered to a mesh of the light source. Their method uses aggregated anisotropic point lights on the boundary surface for the illumination of the scene. Kniep et al. [2009] store photons with directional information on a user-defined boundary surface of a virtual luminaire. Velazquez-Armendariz et al. [2015] pre-compute anisotropic point lights and a radiance volume to accelerate the global illumination of the exterior and the rendering of the luminaire appearance. Krösl et al. [2017] presented a system based on photon tracing and a multi-resolution image filtering technique to generate virtual far-field measurement data of luminaires with complex reflector geometries for interactive prototyping.

3 METHOD OVERVIEW

In this section, we introduce the foundation of our real-time approximation for polygonal photometric lights and discuss our decisions. The reflected radiance L at a point x in view direction ω_o of a photometric light defined in the area \mathcal{A} can be expressed as the following integral over \mathcal{A} :

$$L(x, \omega_o) = \frac{1}{A} \int_{\mathcal{A}} f_r(\mathbf{x}, \omega_i, \omega_o) I(-\omega_i) \frac{(\omega_i \cdot \mathbf{n}_x)}{\|x' - x\|^2} dx', \quad (1)$$

where A is the surface area of \mathcal{A} , f_r is the *Bidirectional Reflectance Distribution Function* (BRDF), ω_i is the incidence direction pointing from x towards x' , I is the function of measured radiant intensity and \mathbf{n}_x is the surface normal of x . We use the normalization of $1/A$ since photometric data represents the absolute emission of a luminaire and the modeled area should not affect the illumination intensity. In this paper, we focus on reflected radiance assuming full visibility. Thus we omit the visibility term, putting our approximation in the same line as other real-time area light shading techniques. Our technique is picking up similar strategies as used with point-based approximations, such as the *Most Representative Point* (MRP) method [Drobot 2014] or *Structured Sampling* [Lagarde and Rousiers 2014]. Unfortunately, simplification is difficult. The function of radiant intensity I is evaluated using ω_i and is defined by measurement data. Also, the definition as an area integral containing $1/\|x' - x\|^2$ makes a robust approximation for all constellation of x and A difficult. To simplify this problem, we reformulate L as an integral over solid angles. This implies that we need to transform the radiant intensity function I depending on the direction ω to the radiance $L_e(\omega)$ emitted from the area A using the following substitution:

$$L_e(\omega) = \frac{I(\omega)}{A(\omega \cdot \mathbf{n}_A)}, \quad (2)$$

where \mathbf{n}_A is the surface normal of the light-emitting area. Note that this does not allow the expression of light transport to points in the plane of the light emitting-area, where $\omega \cdot \mathbf{n}_A = 0$ and potentially introduces numerical issues at glancing angles. We have still chosen this reformulation as I is often 0 or relatively small in these directions and we have found a simple workaround to overcome this limitation (see Section 5.3).

Secondly, we consider a solution for diffuse and specular transport light separately and continue by replacing the general BRDF by a diffuse BRDF $f_r = \frac{\rho}{\pi}$. This allows the BRDF to be moved out of the integral. Together with the radiance substitution (Equation 2), we rewrite Equation 1 for the diffuse reflected radiance L_d as an integral over solid angle Ω subtended by the light source:

$$L_d(x) = \frac{\rho}{\pi} \int_{\Omega} L_e(-\omega_i) \underbrace{(\omega_i \cdot \mathbf{n}_x)}_G d\omega_i \quad (3)$$

We use this equation as basis in our approximation. The geometric term G together with $1/(-\omega_i \cdot \mathbf{n}_A)$ of L_e represents a smoothly changing term over the solid angle Ω subtended by the light source. This increases our expectations to find a practical real-time solution at reasonable costs. In Section 4 we discuss different sampling strategies to approximate Equation 3 and conclude with our technique in Section 5 and its evaluation in Section 6. Finally, we provide a solution for specular reflection in Section 7.

4 SAMPLING STRATEGIES

Given is a photometric light defined in a polygonal region where the measured radiant intensity is emitted from. We have separated the diffuse shading L_d and transformed it to an integral over the solid angle subtended by the region in Equation 3. In this section, we review sampling strategies and simplifications used for diffuse area lights and explore their practicability to approximate this integral.

4.1 Structured Sampling

Lagarde and Rousiers [2014] use structured sampling to approximate direct lighting of diffuse area lights. Instead of using many randomly chosen samples as it is common in Monte Carlo integration, a small number of well-selected samples are used to cover the integration domain, i.e., the area of the light source, sufficiently. The probability density function (PDF) can then be assumed to be constant, reducing the solution to the calculation of an average. Applied to Equation 3 we get:

$$L_d(x) = \frac{\rho}{\pi} \frac{\Omega}{N} \sum_{i=1}^N L_e(-\omega_i) (\omega_i \cdot \mathbf{n}_x). \quad (4)$$

The selection of samples is crucial for the accuracy and visual quality of the approximation. Lagarde and Rousiers use the four corner points and the barycenter of an area light, which yields good results for diffuse emitters. However in our case, with varying L_e , as soon as the emission profile of a light becomes less diffuse, sample patterns become visible. Figure 2 (b) shows this technique used with a data set with high directionality. It can be seen that the samples can be distinguished based on the illumination near the light source. This shows that sampling the corners and the barycenter is not a reasonable choice for arbitrarily emitting area lights.

Although there is room for improvement, such as an increased number of samples, we strive for a more efficient solution. Next, we elaborate on a different sampling strategy.

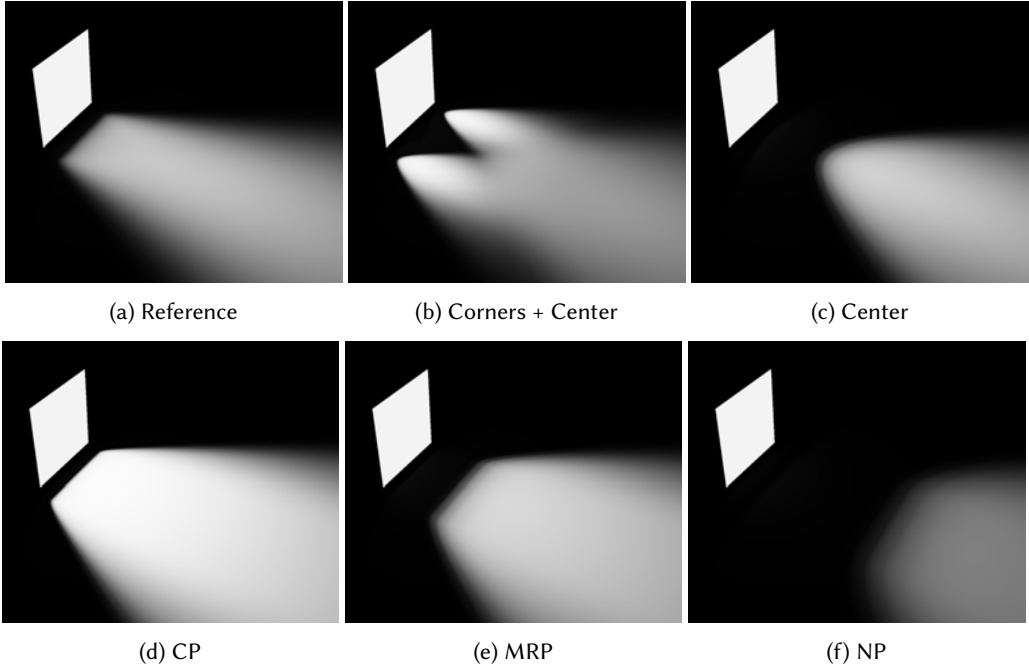


Fig. 2. Comparison of a reference rendering of the ARCOS data set, a luminaire with directed emission (see Section 6) to various approximations, generated by using different simple sampling strategies to calculate the light transport of Equation 3.

4.2 Dynamic Points

Instead of using several well-chosen samples, Drobot [2014] approximates area lights by a single sample. His technique is a form of importance sampling where only the most important sample is used to approximate the whole function. The position of this sample on the light source is referred to as the *Most Representative Point* (MRP). Its position depends on the spatial relation between the illuminated point and the light source. As an exact solution would be too expensive, the *MRP* is typically derived from spatial heuristics. For diffuse area lights only the geometry term of the rendering equation needs to be considered:

$$\frac{(\omega_i \cdot \mathbf{n}_x)}{\|x' - x\|^2} \quad (5)$$

We again use the geometric term from area to point, as it is used by Drobot and suitable for this approximation. The maximum of this term is approximated by finding the maximum of the numerator and the minimum of the denominator respectively and combining the result to get a single point. We refer to the position of the sample which maximizes the numerator as *Normal Point* (NP) and the position of the point which minimizes the denominator as *Closest Point* (CP). The intersection point of the area with the half-vector between these two points form the *MRP*. To generalize this, we refer to such points as *dynamic points*. Our consideration is that these points might be suitable candidates to evaluate Equation 3 and provide a meaningful estimate. Note, in order to allow all points to be suitable candidates we have adapted the construction procedure of Drobot and clamp NP and CP to the nearest point on the light-emitting polygon. The corresponding illustration is shown in Figure 3.

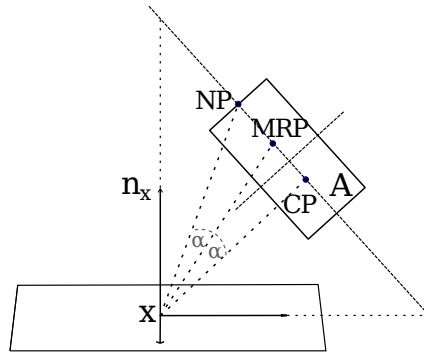


Fig. 3. The 3 dynamic points: normal point NP , closest point CP and the most representative point MRP . The position of the points on the area light A depends on the spatial relation to x .

Using solely any of these points does not generate pleasing results (see Figure 2, bottom), but we can learn about their characteristics. The MRP is in most cases a good average for the overall intensity, while the closest point is most important in defining the shape of the illumination, but overshoots in intensity. Obviously, a sampling strategy covering the entire area of the light source is necessary for the application with photometric lights. Therefore, in our method, we combine *Structured Sampling* with *Dynamic Points*. Next, we complement the sampling strategies with a suitable integration method and elaborate our technique.

5 OUR TECHNIQUE

We propose to combine *Structured Sampling* with *Dynamic Points* as introduced in Section 4. Our experiments showed that these strategies appear to complement each other. While a dynamic point provides sampling pattern-free approximations of area lights, it does not cover the integration domain sufficiently, which, in contrast, is done by structured samples. We incorporate the closest point CP in our sampling strategy, as it is the most important point in defining the shape of the illumination. Finally, we evaluate Equation 3 at the combined sample positions and use a cubature technique based on a triangulation to get suitable sample weights.

5.1 Cubature based on a triangulation

A dynamic point can be anywhere on the light source, therefore it can come close to structured samples or even overlap them. This implies that the Monte Carlo estimator in Equation 4 (structured sampling) cannot be used, because the uniformness of the sample set, which is required to justify the use of the Monte Carlo estimator with a uniform PDF, is lost. In order to approximate the integral of L_d , we adopt a cubature method of scattered data used in Geomathematics [Freedman et al. 2015, p. 1206]. It is based on a triangulation, which provides a partitioning of the integration domain and allows to calculate appropriate sample weights. Given a triangulated sample set D on a sphere the cubature is:

$$Q_N = \sum_{\Delta_{ijk} \in D} \frac{f(x_i) + f(x_j) + f(x_k)}{3} |\Delta_{ijk}| \quad (6)$$

The samples x_i , x_j and x_k do not lie on a common line and are in counterclockwise order as viewed from the center of the sphere (see Figure 4). $|\Delta_{ijk}|$ is the spherical excess of the triangle:

$$|\Delta_{ijk}| = \alpha_i + \alpha_j + \alpha_k - \pi \quad (7)$$

where α_i, α_j and α_k are the angles of Δ_{ijk} .




Fig. 4. Illustration of a spherical triangle as used by the cubature. The vertices are ordered in counterclockwise order relative to the center of the sphere.

Typically, a *Delaunay* triangulation is used to partition the area fairly and to avoid thin long triangles. It is given when the *empty circumcircle interior property* is valid for all triangles [Renka 1997]. For a spherical quadrilateral $\{i, j, k, l\}$ the condition is:

$$\frac{(x_j - x_i) \times (x_k - x_i)}{\|(x_j - x_i) \times (x_k - x_i)\|} \cdot (x_l - x_i) > 0 \tag{8}$$

The resulting weighting scheme makes this integration technique well suited for our combined sample set, as long as we can build a meaningful triangulation. Next, we explain how this integration method is incorporated in our shading procedure.

5.2 Shading Procedure

The first step is to build a triangulation of our integration domain bound by the polygon and include the closest point. This needs to be performed for every shading point x . As only visible parts of the polygon contribute to the shading, we clip the polygon by the horizon plane of x . The next step is to find the closest point CP on the clipped polygon. Therefore, x is projected onto the light's plane and then clamped to the polygon if it is outside. In both cases, fully visible and clipped, there are three positional relationship cases for CP : 1. inside the polygon, 2. clamped to an edge or 3. clamped to a vertex (Figure 5). Clipping will also add an edge with two new vertices, v_a and v_b .




Fig. 5. Illustration of positional relationship cases of the closest point CP (red): 1. inside (left), 2. edge (center) and 3. vertex (right); when the polygon is fully visible (top) and partly clipped (bottom). A valid triangulation is given by building a triangle fan around CP .




Fig. 6. Visualization of violations of the Delaunay condition (red regions) when using our triangulation scheme. In case the polygon is fully visible and not clipped by the horizon of the shaded point the condition is always fulfilled (left). On surfaces where the polygon is clipped and the closest point does not coincide with a corner the triangulation often violates the Delaunay condition (right).

Considering only convex polygons, a simple rule to get a valid triangulation in all cases is to create a triangle fan around CP . Before continuing with the radiance integration, we transform all vertices into tangent space and normalize them to get the vertices of our spherical triangles.

As for our intended integration method an even triangulation is suggested, we have evaluated the compliance of the *Delaunay* condition using Equation 8. Considering the 2d-case in Figure 5 we notice that the condition is fulfilled for all unclipped cases (top). We can also confirm this in a 3d scene for spherical triangles (Figure 6 left). Regions violating the condition (red) are sparsely visible. In clipped cases, we can observe that this condition is often violated when the closest point is on an edge (right). We have experimented with updating the triangulation at runtime and have concluded that this is not worth the computational costs and too complex for general polygons. Therefore, we directly continue with the radiance approximation using a triangle fan around CP .

An accompanying illustration of the following procedure is depicted in Figure 7. Applying the cubature technique of Equation 6 to solve Equation 3, we get:

$$L_d(x) = \frac{\rho}{\pi} \sum_{\Delta_{ijk} \in D} \frac{L_i G_i + L_j G_j + L_k G_k}{3} |\Delta_{ijk}|, \quad (9)$$

where L_i is short for $L_e(x_i)$ introduced in Equation 2 and G_i for $G(x_i)$ in Equation 3. However, we can see that this would cause a sample at the horizon plane ($\mathbf{n}_x \cdot \omega_i = 0$) to get 0 weight. This excludes the closest point and vertices at the horizon from contributing to the sum in certain light constellations. Even close above the horizon the weight is inappropriate and thereby makes this scheme impractical. In order to avoid this and to reduce the impact of potential extrema in L and G we replace their product by the product of their averages:

$$L_d(x) = \frac{\rho}{\pi} \sum_{\Delta_{ijk} \in D} \frac{L_i + L_j + L_k}{3} \frac{G_i + G_j + G_k}{3} |\Delta_{ijk}| \quad (10)$$

This weighted sum over all triangles constitutes the diffuse shading at a point x with a polygonal photometric light. Since the triangulation scheme of our method can generate thin triangles when the closest point moves towards a vertex, the calculation of the spherical excess $|\Delta_{ijk}|$ (Equation 7) needs to gracefully converge to 0 for degenerated triangles. Therefore, we use the vector form [van

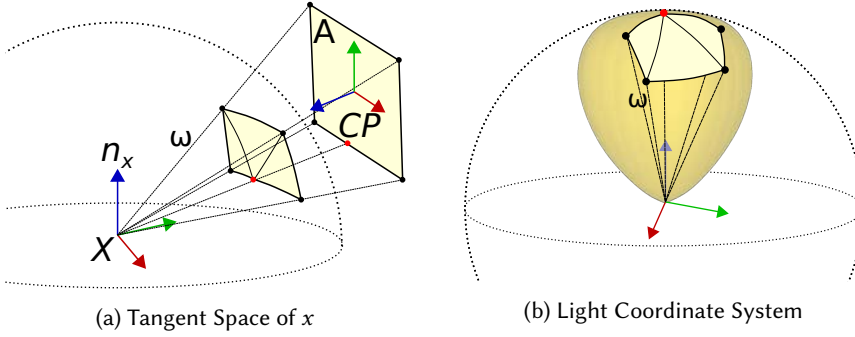


Fig. 7. Illustration of our method. The polygonal light A with constructed closest point CP is projected into tangent space of the point to shade x (a). Radiant intensity is sampled by rotating a sample direction ω into the light coordinate system orientation (b). The triangulation gives the samples weights for the shading calculation.

[Oosterom and Strackee 1983]:

$$\tan\left(\frac{1}{2}\Omega\right) = \frac{\hat{a} \cdot (\hat{b} \times \hat{c})}{1 + (\hat{a} \cdot \hat{b}) + (\hat{a} \cdot \hat{c}) + (\hat{b} \cdot \hat{c})}, \quad (11)$$

where $\Omega = |\Delta_{ijk}|$ and $\hat{a}, \hat{b}, \hat{c}$ are the normalized vectors to the vertices of the spherical triangle. This provides a robust solution in case any of the vectors become identical.

All steps described in this section are executed in a fragment shader. The overall procedure has linear runtime complexity depending on the number of vertices, respectively the number of triangles and edges. For details, we refer to the shader code in the supplemental material. There, the clamping algorithm of CP is limited to convex polygons, as well as our triangulation scheme. The number of vertices is configured with 4 as this is sufficient for our application, but the implementation can handle general polygons.

5.3 Numerical Robustness

We want to elaborate on our countermeasures for the numerical limitations of the radiance substitution (Equation 2). The formulation does not allow us to calculate shading of points within the polygon plane, but photometric data often requires this. In particular, there are two counteracting expressions. The radiance becomes infinite (1) due to the division by $(\omega \cdot \mathbf{n}_A)$ and the solid angle of the polygon converges to zero (2). This also introduces numerical issues close to the polygon horizon. The images of Figure 8 have been rendered with no precautions to prevent numerical issues. In a typical view (a) visual artifacts can be noticed by occasionally popping of individual pixels. In a closeup of the illumination within the polygon horizon (b) the numerically unstable band becomes clearly visible. In the scenario where the light is parallel to a surface and then is moved till both surfaces touch (e.g. moving the light in (a) to the ceiling), the entire surface is susceptible to such artifacts. In order to avoid these artifacts and circumvent the limitation, we shift the shading point away from the horizon at the beginning of our procedure, if it is inside a certain epsilon region. At the same time, we calculate an opposite shift that is applied to the sample directions when evaluating the intensity profile to get a continuous illumination. We decided to skip the opposing shift of the lookup vector in the final implementation as we could not notice any discontinuities after tweaking the size of the epsilon region.




Fig. 8. Images generated without any precautions to prevent numerical issues. In typical conditions (a) almost no artifacts are noticeable. A closeup of the illumination at the horizon shows the manifestation of the singularity introduced by the radiance substitution in Equation 2.

5.4 Summary

Our integration method given in Equation 10 represents a weighted average of the area light's radiance for a point x to calculate the diffuse shading of a polygonal photometric light. Including the closest point CP adds a dynamic point in relation to x and provides a continuous transition when x moves. The size and shape of the integration domain depend on the relation of x and A (see Figure 7). For most positions it will be relatively small and for constellation with large solid angle, we include CP to stabilize the result. Assuming the light has a constant intensity profile, hence it is an omni-directional point light, we only need to solve the integral for $(\mathbf{n}_x \cdot \omega_i)$ by $(-\omega_i \cdot \mathbf{n}_A)$ of G and L_e . Similar to *Structured Sampling* and the *MRP* method, our technique solves this sufficiently (see Section 3 of the supplemental evaluation). When including the intensity profile, we need to transform the lookup-directions ω_i from tangent space of x to the coordinate system of the light source. As this is only a rotation the intensity profile is sampled over an equal shaped region. This allows to conclude possible limitations of our methods. We expect that the approximation is reasonable if the intensity profile's function is relatively smooth over the integration domain. Typical intensity profiles (see Figure 9) indicate that this might be practical. Nevertheless, details within the sampling region will be missed and intensity variations might occur when the samples move over details in their transition. We will continue this discussion during the evaluation in the next section.

6 EVALUATION AND RESULTS

In order to benchmark the accuracy of our real-time approximation technique, we have selected 7 intensity profiles from a luminaire manufacturer [Zumtobel 2019] that each has its own characteristics to challenge our method (see Figure 9). Our evaluation process tests all photometric profiles under normalized conditions. The scene contains a ground plane with normal n_g and an area light of 1×1 units at heights of $\{0.1, 1.0, 2.5\}$ units above the ground and orientations of $\cos^{-1}(n_A \cdot n_g) \in \{\pi/2, 5/8\pi, 3/4\pi, 7/8\pi, \pi\}$. We render images from a top down view with a field of view covering a region of 5 units around the light source. We apply tone mapping with $L_d(x, y) = L(x, y)/(1 + L(x, y))$ of $L(x, y) = (0.05/\overline{L_w}) * \overline{L_w}$, where $\overline{L_w}$ is the logarithmic average luminance. The reference image (row 3) is rendered with the method of Urena et al. [2013] using 20,000 samples. The next row shows our approximation *Cubature* (row 4) and the difference images at height 0.1 and orientation $\pi/2$ (row 5). Our method always provides well-defined shapes of the




Fig. 9. Intensity profiles of the luminaires we use in our evaluation (row 1) [Zumtobel 2019]. Screenshots from our showroom scene with luminaire models linked to the selected data sets (row 2). The reference renderings from our evaluation setup with $h = 0.1$ and orientation $\pi/2$ (row 3). Approximation images rendered using our *Cubature* method (row 4). False color difference images (row 5) where red areas are too bright and blue areas are too dark. The columns are labeled with a short name of the data set.

illumination that all appear convincing without a reference image. Looking at the difference images reveals certain tendencies of our method. It works very well for almost diffuse intensity profiles (SLOTLIGHT & PERLUCE). Looking at more directed profiles (ARCOS & PANOS), the light cones appear too wide. We suspect this is caused by the dominance of the closest point. We also notice that there are some structures in regions that should be smooth (MIREL & LINETIK). In these cases they are caused by a rectangular cross-section of the intensity profiles when viewed from the polar axis. Our approximation reproduces this variation in the near-field where the dynamic point moves along the edge of the polygon. Likewise, in cases where the intensity profiles has fine details or tiny peaks (INTRO) they are emphasized due the moving point, but the rendering still appears plausible and produces a robust approximation of the remaining illumination. In practice, most intensity profiles are rather smooth as luminaires are measured in large distances where caustics are typically absent and luminaire engineers often optimize for uniformity. Another fact that accommodates our technique is that the near-field is typically the brightest region in an image and highly compressed by tone mapping. Textured surfaces or variations of the surface normal also make it harder to recognise intensity variations as artifacts (see Figure 12). We found that our approximation is well suited for real-world data sets. We refer to the supplemental material for the evaluation results of all configurations and to the accompanying video for a demonstration.

The interactive segment renders 16 photometric light simultaneously and includes point-based shadows using cube shadow maps.

Additionally, we have quantified the difference using the NRMS error metric. Figure 10 shows the combined result for each data set per height, calculated from stitched images containing all orientations. The error correlates with the difference images, but does not directly imply the plausibility of the generated images. An example are the data sets with directed intensity profiles of ARCOS and PANOS, the visual appearance suggest a very good approximation, but is actually not that accurate due to the slightly distorted shape of the outline.

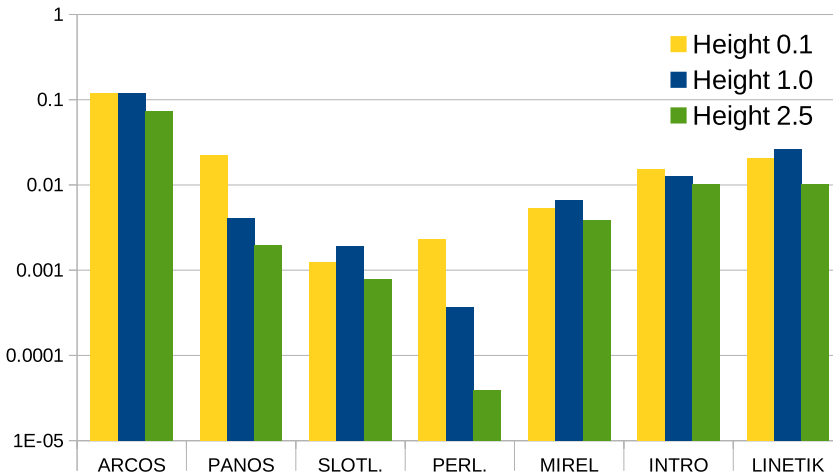


Fig. 10. NRMS error calculated from stitched images showing the illumination of all 5 different light orientations at polygon heights of 0.1, 1.0 and 2.5 units.

Finally, we evaluate the performance of our method. Table 1 lists the frame times in HD and UHD resolution, with diffuse shading, no visibility test, tone mapping disabled and a 4-channel 32bit floating point frame buffer, measured on two different hardware systems. Our method is bound by its computations in the fragment shader and scales linearly with the resolution. The baseline is given by the *Point* approximation, performing a single lookup on the intensity profile and only calculating the distance attenuation to the centroid of the polygon. Our method (*Cubature*) takes about four times longer while using 4-5 samples in typical cases. The reference implementation (*MC 8*) is able to perform Monte Carlo integration with 8 random samples slightly faster and would allow about 10 samples at equal time. The combination with a 16 pixel radius separated Gaussian de-noising filter (*DN 2x16*) is much more expensive. A visual comparison of the Monte Carlo

Table 1. Frame times of our method (bold) and others.

Method	NV RTX 2080 Ti		NV GTX 1070 Ti	
	1080p	2160p	1080p	2160p
Cubature	0.41 ms	1.70 ms	1.12 ms	4.56 ms
Cubature + LTC	0.77 ms	3.15 ms	1.81 ms	7.38 ms
Point	0.11 ms	0.40 ms	0.17 ms	0.67 ms
MC 8	0.33 ms	1.41 ms	0.89 ms	3.51 ms
MC 8 + DN 2x16	0.94 ms	3.02 ms	2.17 ms	9.30 ms

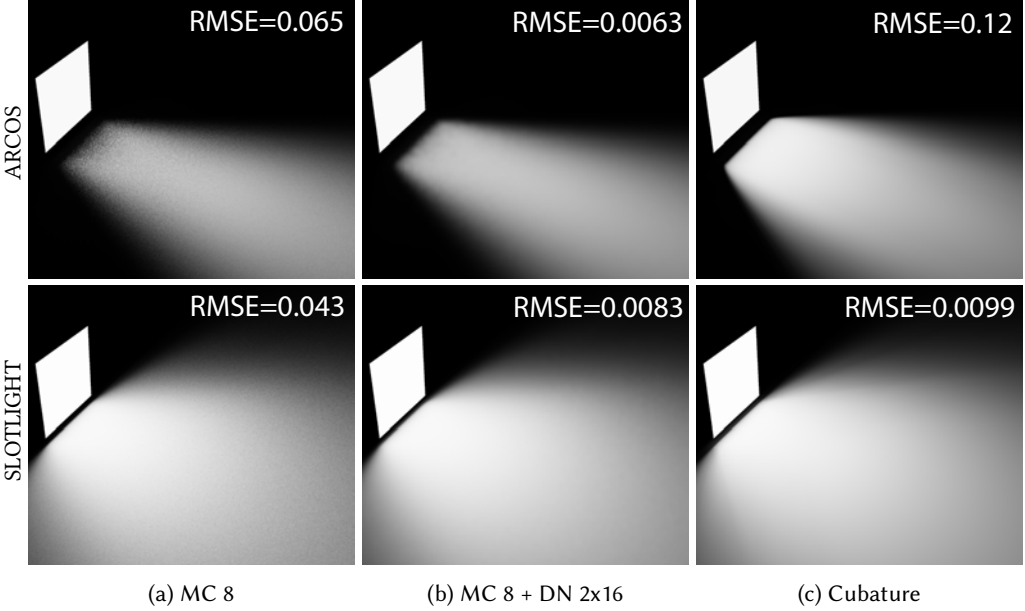


Fig. 11. Comparison of a Monte Carlo solution with 8 samples (*MC 8*) and its de-noised image (*MC 8 + DN 2x16*) to our method (*Cubature*).

solution and our method is shown in Figure 11. Please note that our de-noising method is intended to provide a performance reference and does not represent any state-of-the-art implementation. We have selected the ARCOS and SLOTLIGHT data sets where our method has the highest and lowest error. The MC solution is able to provide a consistent level of quality over all data sets, but de-noising artifacts can be visible. In comparison, our technique produces a noise-free and similarly convincing illumination in shorter time. Our extension for specular reflections (*Cubature + LTC*, see Section 7) is also well suited for real-time applications.

7 SPECULAR SHADING EXTENSION

In order to provide a full shading solution for physically based rendering systems, we propose an extension for the specular reflected radiance L_s . In the extreme case, when the BRDF is a perfect reflection, we can directly evaluate $L_e(\omega)$ using the reflected view direction. In other cases generalized by a microfacet BRDF with arbitrary roughness, we need to calculate a convolution with the BRDF over a certain region or the entire area in the other extreme. We simplify this general case by assuming the extreme, and calculate an energy-equivalent radiance \bar{L}_e to a diffuse area light and use this as supplement for specular shading. We then perform the shading calculation with the *Linearly Transformed Cosines* (LTC) [Heitz et al. 2016] method:

$$L_s(x, \omega_o) = \bar{L}_e \cdot LTC(x, \omega_o, A), \quad (12)$$

where the function *LTC* approximates the convolution of the specular part of the BRDF f_r and the area light A represented as polygon. The constant factor \bar{L}_e can be derived by setting the expression of Equation 3 equal to the light transport of a diffuse light:

$$\int_{\Omega} L_e(-\omega_i)(\omega_i \cdot \mathbf{n}_x) d\omega_i = \bar{L}_e \int_{\Omega} (\omega_i \cdot \mathbf{n}_x) d\omega_i, \quad (13)$$




Fig. 12. Real-time rendering of photometric lights with our method and specular shading extension. The shadows are rendered point-based using cube shadow maps.

and solving it for $\overline{L_e}$:

$$\overline{L_e} = \frac{\int_{\Omega} L_e(-\omega_i)(\omega_i \cdot \mathbf{n}_x) d\omega_i}{\underbrace{\int_{\Omega} (\omega_i \cdot \mathbf{n}_x) d\omega_i}_{G_{\Omega}}} \quad (14)$$

The resulting equation for $\overline{L_e}$ means that we need to normalize the integral for the diffuse illumination of Equation 3 by the integral of the geometric term G , that we call G_{Ω} . This integral term is well studied and real-time approximations, such as *Structured Sampling* or the *MRP* method, provide reasonable approximations. Likewise, our cubature technique achieves similar accuracy for diffuse emission (see Section 3 of the supplemental evaluation). This allows us to base the calculation of $\overline{L_e}$ on our cubature method and approximate G_{Ω} alongside with $L_d(x)$ (Equation 10) at almost no additional costs:

$$G_{\Omega} = \sum_{\Delta_{ijk} \in D} \frac{G_i + G_j + G_k}{3} |\Delta_{ijk}| \quad (15)$$

We then use G_{Ω} according to Equation 14 and calculate the energy-equivalent radiance $\overline{L_e}$ to a diffuse area light:

$$\overline{L_e} = \frac{L_d(x)}{G_{\Omega}(dx)} \quad (16)$$

The energy-equivalent radiance $\overline{L_e}$ is then used to calculate the specular reflection by applying the LTC shading technique (Equation 12). Figure 12 shows two renderings of a scene shaded from multiple photometric polygonal lights using our method and this extension.

8 CONCLUSIONS AND FUTURE WORK

Our real-time approximation technique of photometric polygonal lights provides a valuable extension to real-time rendering systems. Our cubature generates noise-free illumination in the near-field of area lights at reasonable costs and closes the gap to stochastic methods. We have shown its practicability on a variety of data sets and achieved results closely resembling the reference renderings. Our extensions for specular shading provide a complete solution for physically based rendering

facilitating existing techniques for diffuse area lights. We see potential in numerous applications, particularly in architecture visualization and lighting design.

In our evaluation we could identify artifacts due to high-frequency details in the measurement data. This issue could be addressed by introducing a pre-processing step for the intensity profile to provide pre-integrated data at runtime. For example, a hierarchy of smoothed representations could provide a robust average, depending on the shading point, and reduce banding artifacts like the ones visible in the false color images in Figure 9. Additionally, the limitation to polygonal light shapes suggest to extend our method with specialized solution for linear and disk shaped lights. Adding a fixed or second dynamic point to our cubature technique can increase accuracy at minimal additional costs. However, for applications with a high demand on accuracy, elaborated offline rendering techniques might be more suitable. In this field, it would be particularly interesting to evaluate the applicability of *Product Importance Sampling* techniques for photometric lights.

ACKNOWLEDGMENTS

We thank our colleges Harald Steinlecher and Andreas Walch for their insightful and valuable comments on this work. We thank Bert Junghans, Christian Bauer and Sammy Suleiman-Fröhlich from Zumtobel for the ongoing fruitful collaboration and for providing test data used to evaluate this work.

VRVis is funded by BMVIT, BMDW, Styria, SFG and Vienna Business Agency in the scope of COMET - Competence Centers for Excellent Technologies (854174) which is managed by FFG.

REFERENCES

- James Arvo. 1995. Applications of Irradiance Tensors to the Simulation of non-Lambertian Phenomena. In *Proceedings of the 22Nd Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '95)*. ACM, New York, NY, USA, 335–342.
- I. Ashdown. 1993. Near-Field Photometry: A New Approach. *Journal of the Illuminating Engineering Society* 22, 1 (1993), 163–180.
- Ian Ashdown. 1995. Near-Field Photometry: Measuring and Modeling Complex 3-D Light Sources. In *SIGGRAPH 1995*.
- I. Ashdown, L. Bedocs, W. Carroll, J. Boer, P. Dehoff, Michael Donn, H. Erhorn, L. Escaffre, Marc Fontoyntont, and Phillip Greenup. 2006. *CIE 171:2006 Test Cases to Assess the Accuracy of Lighting Computer Programs*.
- Ian Ashdown and Ron Rykowski. 1998. Making Near-Field Photometry Practical. *Journal of the Illuminating Engineering Society* 27, 1 (1998), 67–79.
- Daniel R. Baum, Holly. E. Rushmeier, and James M. Winget. 1989. Improving Radiosity Solutions Through the Use of Analytically Determined Form-Factors. In *Proceedings of the 16th annual conference on Computer graphics and interactive techniques*. ACM, New York, USA, 325–334.
- Nikolaus Binder, Sascha Fricke, and Alexander Keller. 2019. Massively Parallel Path Space Filtering. *CoRR* abs/1902.05942 (2019). arXiv:1902.05942
- David Burke, Abhijeet Ghosh, and Wolfgang Heidrich. 2005. Bidirectional Importance Sampling for Direct Illumination. In *Proceedings of the Sixteenth Eurographics Conference on Rendering Techniques (EGSR '05)*. Eurographics Association, Aire-la-Ville, Switzerland, Switzerland, 147–156.
- Petrik Clarberg, Wojciech Jarosz, Tomas Akenine-Möller, and Henrik Wann Jensen. 2005. Wavelet Importance Sampling: Efficiently Evaluating Products of Complex Functions. *ACM Trans. Graph.* 24, 3 (July 2005), 1166–1175.
- Robert L. Cook, Thomas Porter, and Loren Carpenter. 1984. Distributed Ray Tracing. *SIGGRAPH Comput. Graph.* 18, 3 (Jan. 1984), 137–145.
- Carsten Dachsbacher, Jaroslav Krivánek, Miloš Hašan, Adam Arbree, Bruce Walter, and Jan Novák. 2014. Scalable Realistic Rendering with Many-Light Methods. *Comput. Graph. Forum* 33, 1 (Feb. 2014), 88–104.
- David L DiLaura, Kevin W Houser, Richard G Mistrick, and Gary R Steffy. 2011. *The Lighting Handbook: Reference and Application*. Illuminating Engineering Society of North America New York (NY).
- Pierre Yves Donzallaz. 2019. *Create High-Quality Light Fixtures in Unity*. Technical Report. Unity Technologies.
- Míchal Drobot. 2014. *GPU Pro 5*. A K Peters/CRC Press, Chapter Physically Based Area Lights, 67–100.
- Alejandro Conty Estevez and Pascal Lecocq. 2018. Fast Product Importance Sampling of Environment Maps. In *ACM SIGGRAPH 2018 Talks (SIGGRAPH '18)*. ACM, New York, NY, USA, Article 69, 2 pages.

- Willi Freeden, M. Zuhair Nashed, and Thomas Sonar (Eds.). 2015. *Handbook of Geomathematics* (2 ed.). Springer-Verlag, Berlin Heidelberg.
- Cindy M. Goral, Kenneth E. Torrance, Donald P. Greenberg, and Bennett Battaile. 1984. Modeling the Interaction of Light Between Diffuse Surfaces. *SIGGRAPH Comput. Graph.* 18, 3 (Jan. 1984), 213–222.
- Ibón Guillén, Carlos Ureña, Alan King, Marcos Fajardo, Iliyan Georgiev, Jorge López-Moreno, and Adrian Jarabo. 2017. Area-Preserving Parameterizations for Spherical Ellipses. *Comput. Graph. Forum* 36, 4 (July 2017), 179–187.
- Peter Hedman, Tero Karras, and Jaakko Lehtinen. 2016. Sequential Monte Carlo Instant Radiosity. In *Proceedings of the ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games*. ACM.
- Wolfgang Heidrich, Jan Kautz, Philipp Slusallek, and Hans-Peter Seidel. 1998. Canned Lightsources. In *Rendering Techniques '98*, George Drettakis and Nelson Max (Eds.). Springer Vienna, Vienna, 293–300.
- Eric Heitz, Jonathan Dupuy, Stephen Hill, and David Neubelt. 2016. Real-time Polygonal-light Shading with Linearly Transformed Cosines. *ACM Trans. Graph.* 35, 4, Article 41 (July 2016), 8 pages.
- Eric Heitz and Stephen Hill. 2017. *GPU Zen: Advanced Rendering Techniques*. Chapter Linear-light shading with linearly transformed cosine.
- Eric Heitz, Stephen Hill, and Morgan McGuire. 2018. Combining Analytic Direct Illumination and Stochastic Shadows. In *ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games*. 10. I3D 2018.
- James T. Kajiya. 1986. The Rendering Equation. In *Proceedings of the 13th Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '86)*. ACM, New York, NY, USA, 143–150.
- Brian Karis. 2013. Real Shading in Unreal Engine 4. *part of ACM SIGGRAPH 2013 Course: Physically Based Shading in Theory and Practice* (2013).
- Alexander Keller. 1997. Instant Radiosity. In *Proceedings of the 24th annual conference on Computer graphics and interactive techniques (SIGGRAPH '97)*. ACM Press/Addison-Wesley Publishing Co., New York, NY, USA, 49–56.
- Stefan Knip, S. Häring, and Marcus Magnor. 2009. Efficient and Accurate Rendering of Complex Light Sources. *Computer Graphics Forum* 28 (08 2009), 1073 – 1081.
- Katharina Krösl, Christian Luksch, Michael Schwärzler, and Michael Wimmer. 2017. LiteMaker: Interactive Luminaire Development using Progressive Photon Tracing and Multi-Resolution Upsampling. In *Vision, Modeling & Visualization*, Matthias Hullin, Reinhard Klein, Thomas Schultz, and Angela Yao (Eds.). The Eurographics Association, 1–8.
- S. Lagarde and C.D. Rousiers. 2014. Moving Frostbite to Physically Based Rendering. *part of ACM SIGGRAPH2014 Course: Physically Based Shading in Theory and Practice* (2014).
- Gilles Laurent, Cyril Delalandre, Grégoire de La Rivière, and Tamy Boubekeur. 2016. Forward Light Cuts: A Scalable Approach to Real-Time Global Illumination. *Comput. Graph. Forum* 35, 4 (July 2016), 79–88.
- Pascal Lecocq, Arthur Dufay, Gaël Sourimant, and Jean-Eudes Marvie. 2016. Accurate Analytic Approximations for Real-time Specular Area Lighting. In *Proceedings of the 20th ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games (I3D '16)*. ACM, New York, NY, USA, 113–120.
- Daqi Lin and Cem Yuksel. 2019. Real-Time Rendering with Lighting Grid Hierarchy. *Proc. ACM Comput. Graph. Interact. Tech. (Proceedings of I3D 2019)* 2, 1, Article 8 (2019), 17 pages.
- Christian Luksch, Robert F. Tobler, Ralf Habel, Michael Schwärzler, and Michael Wimmer. 2013. Fast Light-Map Computation with Virtual Polygon Lights. In *Proceedings of ACM Symposium on Interactive 3D Graphics and Games 2013*. ACM, New York, NY, USA, 87–94.
- Christian Luksch, Michael Wimmer, and Michael Schwärzler. 2019. Incrementally Baked Global Illumination. In *Proceedings of the 33rd Symposium on Interactive 3D Graphics and Games (I3D '19)*, Ari Rapkin Blenkhorn (Ed.). ACM.
- Michael Mara, Morgan McGuire, Benedikt Bitterli, and Wojciech Jarosz. 2017. An Efficient Denoising Algorithm for Global Illumination. In *Proceedings of High Performance Graphics*. ACM, New York, NY, USA.
- Albert Mas, Ignacio Martín, and Gustavo Patow. 2008. Compression and Importance Sampling of Near-Field Light Sources. *Computer Graphics Forum* 27, 8 (2008), 2013–2027.
- Morgan McGuire. 2010. Ambient Occlusion Volumes. In *Proceedings of the 2010 ACM SIGGRAPH Symposium on Interactive 3D Graphics and Games (I3D '10)*. ACM, New York, NY, USA, Article 12, 1 pages.
- Jacob Munkberg, Jon Hasselgren, Petrik Clarberg, Magnus Andersson, and Tomas Akenine-Möller. 2016. Texture Space Caching and Reconstruction for Ray Tracing. *ACM Trans. Graph.* 35, 6, Article 249 (Nov. 2016), 13 pages.
- Christoph Peters and Carsten Dachsbacher. 2019. Sampling Projected Spherical Caps in Real Time. *Proc. ACM Comput. Graph. Interact. Tech.* 2, 1, Article 1 (June 2019), 16 pages.
- Robert J. Renka. 1997. Algorithm 772: STRIPACK: Delaunay Triangulation and Voronoi Diagram on the Surface of a Sphere. *ACM Trans. Math. Softw.* 23, 3 (Sept. 1997), 416–434.
- Christoph Schied, Anton Kaplanyan, Chris Wyman, Anjul Patney, Chakravarty R. Alla Chaitanya, John Burgess, Shiqiu Liu, Carsten Dachsbacher, Aaron Lefohn, and Marco Salvi. 2017. Spatiotemporal Variance-guided Filtering: Real-time Reconstruction for Path-traced Global Illumination. In *Proceedings of High Performance Graphics (HPG '17)*. ACM, New York, NY, USA, Article 2, 12 pages.

- Peter Shirley, Changyaw Wang, and Kurt Zimmerman. 1996. Monte Carlo Techniques for Direct Lighting Calculations. *ACM Trans. Graph.* 15, 1 (Jan. 1996), 1–36.
- Carlos Ureña and Iliyan Georgiev. 2018. Stratified Sampling of Projected Spherical Caps. *Comput. Graph. Forum* 37 (2018), 13–20.
- Carlos Ureña, Marcos Fajardo, and Alan King. 2013. An Area-Preserving Parametrization for Spherical Rectangles. *Computer Graphics Forum* (2013).
- Adriaan van Oosterom and Jan Strackee. 1983. The Solid Angle of a Plane Triangle. *IEEE Transactions on Biomedical Engineering* BME-30 (1983), 125–126.
- Edgar Velázquez-Armendáriz, Zhao Dong, Bruce Walter, and Donald P. Greenberg. 2015. Complex Luminaires: Illumination and Appearance Rendering. *ACM Trans. Graph.* 34, 3, Article 26 (May 2015), 15 pages.
- Channing Verbeck and Donald Greenberg. 1984. A Comprehensive Light-Source Description for Computer Graphics. *IEEE Comput. Graph. Appl.* 4, 7 (July 1984), 66–75.
- Lifeng Wang, Zhouchen Lin, Wenle Wang, and Kai Fu. 2008. *One-Shot Approximate Local Shading*. Technical Report. Tech. rep.
- Gregory J. Ward. 1994. The RADIANCE Lighting Simulation and Rendering System. In *Proceedings of the 21st Annual Conference on Computer Graphics and Interactive Techniques (SIGGRAPH '94)*. ACM, New York, NY, USA, 459–472.
- Zumtobel. 2019. Products - Zumtobel. <https://www.zumtobel.com/com-en/products.html>.