

Computing Exact Shadow Irradiance Using Splines

Michael M. Stark
University of Utah

Elaine Cohen
University of Utah

Tom Lyche
University of Oslo

Richard F. Riesenfeld
University of Utah

Abstract

We present a solution to the general problem of characterizing shadows in scenes involving a uniform polygonal area emitter and a polygonal occluder in arbitrary position by manifesting shadow irradiance as a spline function. Studying generalized prism-like constructions generated by the emitter and the occluder in a four-dimensional (shadow) space reveals a simpler intrinsic structure of the shadow as compared to the more complicated 2D projection onto a receiver. A closed form expression for the spline shadow irradiance function is derived by twice applying Stokes' theorem to reduce an evaluation over a 4D domain to an explicit formula involving only 2D faces on the receiver, derived from the scene geometry. This leads to a straightforward computational algorithm and an interactive implementation. Moreover, this approach can be extended to scenes involving multiple emitters and occluders, as well as curved emitters, occluders, and receivers. Spline functions are constructed from these prism-like objects. We call them generalized polyhedral splines because they extend the classical polyhedral splines to include curved boundaries and a density function. The approach can be applied to more general problems such as some of those occurring in radiosity, and other related topics.

CR Categories: I.3.0 [Computer Graphics]: General; I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism;

Keywords: rendering, shadow algorithms, illumination, visibility determination

1 INTRODUCTION

While computing shadows from a single point source is relatively simple and goes back to the early days of graphics when the light source was moved away from the eye position, the accurate representation of shadows resulting from area sources is quite complicated, as Soler and Sillion [38] point out:

The calculation of detailed shadows remains one of the most difficult challenges in computer graphics, especially in the case of extended (linear or area) light sources.

University of Utah Department of Computer Science, 50 S. Central Campus Dr., 3190 MEB, Salt Lake City, UT 84112

Institutt for informatikk, Universitetet i Oslo, Postboks 1080 Blindern, 0316 Oslo

Email: {mstark,cohen,rfr}@cs.utah.edu, tom@ifi.uio.no

In this paper we present a solution to the general differential area to polygon form factor in the presence of a polygonal occluder, by representing the induced irradiance as a spline function. We generalize the polyhedral spline [4, 16, 32] to obtain our result. The polyhedral spline is defined as a distribution. A geometric interpretation of this definition involves regarding the splines as cross-sectional volumes of higher dimensional polyhedra [14]. The more familiar univariate and tensor product B-splines are special cases of the polyhedral spline.

To introduce the fundamental ideas of our approach, we begin with an analysis in Flatland [26], where the higher dimension construction can be directly visualized and more easily understood. Then we work through the construction of a generalized perspective prism in 4D associated with an element that can be either an emitter or an occluder. The 4D perspective prism is defined so that the 2D cross-section taken through any point (x, y) is a perspective view of the element from the point $(x, y, 0)$. A shadow function can be built by considering boolean operations on such prisms. In the case where the emitter E , occluder O , and receiver R are all parallel, the intersection region of the E and O prisms is isomorphic to the 4D Cartesian product $O \times E$, where O and E are taken as 2D polygons. To bolster confidence for looking into 4D, we observe that the projection of $O \times E$ onto R generates all the discontinuity lines of the shadow structure.

The boolean combination of the O - and E -prisms gives a polyhedron-like object in 4D which is used to define a polyhedral spline function. At each point (x, y) in the plane the value of the spline is obtained by integrating a density function over the 2D cross-section at (x, y) . With a proper definition of the density function this spline gives us the exact partially occluded irradiance at (x, y) . Using Stokes' theorem we derive a recurrence relation that yields a closed form solution for the irradiance function.

1.1 Related Work

Accurate computation of shadows has been a rich and active area of graphics research since near the beginning of the field. In 1990 Woo *et al.* [44] made an extensive survey of relevant techniques. Since then many techniques have been aimed at faster or more accurate computation.

Methods that are based on a single point light source create hard shadows. A typical example of such a method is used in ray tracing, where a ray is cast between a point lying on the surface and a specified light source. The algorithm must determine if there is an occluder between it and the light source [42]. Soft shadows caused by area light sources are simulated by stochastically selecting a set of sample points on the source and casting several rays towards them from the surface [11, 36]. Exact evaluation can be obtained by clipping all the occluders against the area source [18, 33].

Several methods consider the view from a point light source. The shadow buffer creates a discrete depth image from the point light source view [43], using only a single point. Shadow volumes are constructed from a point light source and polygonal occluders [13] or curved boundary occluders [27], from which a point can be tested for its visibility in object space. Soft shadows can be approximated in a similar manner to the ray tracing method, by breaking the light

source into several point samples [7].

The radiosity method theoretically includes the computation of soft shadows. However, the method requires the reconstruction of irradiance functions, for which no general closed form is known. Methods to compute radiosity use a mesh of finite surface elements, then “form factors” representing the exchange of energy between mesh elements are computed. Frequently, the implementation assumes a constant radiosity distribution across each surface patch and finite element, then the values are interpolated across each mesh element. Meshing surfaces without considering the shadow structure can lead to inaccurate solutions and “shadow leaks”. The discontinuity meshing approach [17, 25, 30] was developed to avoid these problems. For each source, a discontinuity mesh is created on the scenes surfaces (receivers). The mesh includes all lines where the illumination function has discontinuities, including the effects of boundaries of multiple occluders. Meshing concerns have been a principal driving force behind shadow research.

It was recognized quite early that solutions could be improved if the form factors could be computed analytically. Analytical solutions have been determined for specific geometries [28, 35]. Generally such solutions assume that each surface is flat and fully visible from the other. Baum *et al* [3] proposed computing the unoccluded point to polygon form factor analytically. To avoid violation of visibility assumptions, they proposed that each source patch be subdivided until all components were either fully visible or fully hidden from all elements in the environment.

Many current methods to compute form factors approximate them with an upper bound calculation, that is, they assume that an unoccluded radiosity kernel can be multiplied against a visibility factor, which gives the fraction of the area of the light source element visible from the receiver [9, 24, 40]. The point to polygon form factor is assumed to have constant exitance on the source. Then it is separated into the unoccluded point to polygon form factor integral and a visibility factor integral.

Arvo [1] provided hope that a closed-form formula for shadow irradiance might exist by developing a clean, elegant closed-form formula for the irradiance Jacobian due to a partially occluded polygonal source. Recently Soler and Sillion [38] showed that when polygonal emitter, occluder, and receiver are all parallel to each other, shadow computations on the receiver can be written exactly as a convolution. This method uses the characteristic functions of both the emitter and the occluder, and computationally uses fast Fourier transforms to compute convolutions. When the elements are not parallel to each other, the method requires approximating the elements by parallel elements, called virtual geometry. Since the error can be estimated, it is suggested that when the error is too large, the elements should be subdivided and new orientations and approximating parallel sub-elements should be chosen until the error is within bounds. They compute irradiance using the radiosity kernel-visibility factor approximation.

Polyhedral splines have been well studied for special polyhedra like simplices, boxes, and cones; de Boor and Höllig’s original paper [4] contains references and a recurrence relation valid for more general polyhedra than simplices, boxes, and cones. Simplex and box splines have been used in geometric modeling [19]. McCool has used polyhedral splines for signal processing and rendering analytically antialiased polygons, and also suggested they might be used for penumbral approximation [31, 32].

2 FLATLAND

We begin by carrying through our formulation in Edwin Abbott’s “Flatland.” In Flatland, surfaces become plane curves, planes become lines, and polygons become line segments [26]. Irradiance is




Figure 1: A general irradiance situation in Flatland. To compute the irradiance at $\mathbf{r}(t)$, we transform to the origin and rotate so that the surface normal is $(0, 1)$. Then the transformed emitter is projected onto the line $y = 1$, where it is a line segment projectively equivalent to the apparent silhouette of $T(E)$. The transformation and resulting objects are all functions of t .

computed from the integral

$$I(r) = \int_C L(r, \theta) \cos \theta d\theta \quad (1)$$

where $L(r, \theta)$ is the radiance coming into r from direction θ , and C is the open semi-circle above the surface at r .

2.1 Formulation of Irradiance

Consider a receiver surface R described by a plane parametric curve $\mathbf{r}(t) = (x(t), y(t))$ and a normal vector $\mathbf{n}(t) = (n_x(t), n_y(t))$. Let E be a uniformly emitting surface, situated above the receiver R , as shown in Figure 1. In order to compute the irradiance at a point $\mathbf{r}(t_0) = (x_0, y_0)$ on R corresponding to a parameter value $t = t_0$ we proceed as follows:

1. Transform the scene such that (x_0, y_0) is mapped to the origin, and such that the unit normal $\mathbf{n}(t_0)$ becomes parallel to the y -axis. This transformation, which we call T , is a translation followed by a rotation, which transforms E into a new position $T(E)$.
2. Let E^* be the perspective projection of $T(E)$ onto the line $y = 1$ with the origin as the center of projection.

We then have

$$E^* = \{(x, 1) : p_1^*(t) \leq x \leq p_2^*(t)\}$$

where

$$p_i^*(t) = \frac{w_{i,1}(t)}{w_{i,2}(t)}, \quad i = 1, 2,$$

and $w_i(t) = (w_{i,1}(t), w_{i,2}(t))$ are the silhouette points of $T(E)$ as seen from the origin. The object $T(E)$ and the projection E^* produce the same irradiance at the origin, which is the irradiance at $\mathbf{r}(t)$. Denoting this irradiance $I_E(\mathbf{r}(t))$ by $I(t)$, we have

$$I(t) = K \left| \int_{p_1^*(t)}^{p_2^*(t)} \frac{\cos \theta \cos \theta'}{d(x)} dx \right|,$$

where K is a constant, $d(x)$ the distance between the origin and the point $(x, 1)$, θ the angle between the ray $(0, 0) \rightarrow (x, 1)$ and the y -axis, and θ' the angle between the same ray and the normal to E^* .

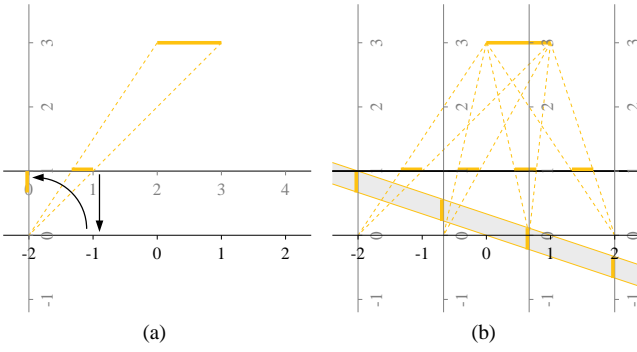


Figure 2: Construction of the strip for a parallel polygon in Flatland. (a) For irradiance at $x = -2$, E is projected to E^* through $x = -2$, then E^* is rotated to the slice of the M_E prism; (b) The slices as a function of x sweep out M_E .

Since E^* is normal to the y -axis we have $\theta' = \theta$, $\cos \theta = 1/d(x)$, and $d(x)^2 = 1 + x^2$. It follows that

$$I(t) = K \left| \int_{p_1^*(t)}^{p_2^*(t)} \frac{dx}{(1+x^2)^{\frac{3}{2}}} \right| = K |\Omega(p_2^*(t)) - \Omega(p_1^*(t))|, \quad (2)$$

where

$$\Omega(x) = \frac{x}{\sqrt{1+x^2}}.$$

The function $I(t)$ can be interpreted geometrically. Suppose for simplicity that $p_2^*(t) \geq p_1^*(t)$ for all t . We consider the integral

$$\int_{p_1^*(t)}^{p_2^*(t)} \omega(y) dy,$$

where

$$\omega(y) = \frac{1}{(1+y^2)^{\frac{3}{2}}}.$$

Let M_E be the 2-dimensional strip

$$M_E = \{(t, y) : -\infty < t < \infty, p_1^*(t) \leq y \leq p_2^*(t)\}.$$

For fixed t the integral is obtained by taking a vertical cut through M_E at t and integrating the density function from the lower boundary of M_E to the upper boundary of M_E . Thus $I(t)$ is the weighted length of this cut at t .

2.2 Irradiance in the Presence of Occluders

Suppose that we extend the scene above by also including an occluder O . In the same way as we did for E we obtain an infinite strip M_O corresponding to the occluder. Where the occluder strip overlaps the emitter strip, the covered portion of the emitter is not visible. The occluded irradiance can be computed from the integral of the same density function $\omega(y)$ over the uncovered portion of the strip. This portion of the strip is the set difference $M_E - M_O$.

If the set difference $M_E - M_O$ is endowed with the density function $\omega(y)$, the integral over each slice of the strip produces the unoccluded irradiance at that point, as in Figure 4(b). This approach generalizes. Suppose we have n occluders O_1, \dots, O_n situated below E . The irradiance at $r(t)$ can be found by integrating the same density function over the set

$$M_E - M_{O_1} - \dots - M_{O_n}.$$

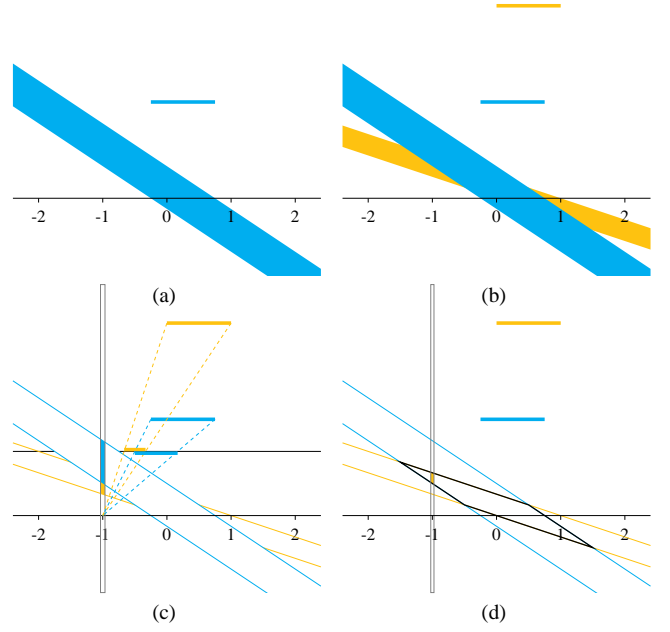


Figure 3: Any Flatland polygon parallel to the floor generates a strip. (a) A polygon closer to the floor generates a thicker and more steeply sloped strip. (b) The cyan strip covers the yellow strip in a way that encodes the occlusion of the yellow emitting polygon by the occluding cyan polygon. (c) A vertical slice of the two strips (shown at $x = -1$) contains a snapshot of the scene. The uncovered portion of the yellow slice corresponds to the unoccluded portion of the emitter. (d) Slices of the intersection of the strips (the black parallelogram) give the occluded portion of the emitter.

If there are also multiple emitters, the total irradiance is computed by summing the contributions from the individual emitters; however, all other emitters must be treated as potential occluders, and there is also the issue of front-and-back culling.

The advantage of this construction is that all the local visibility calculations can be replaced by a single set operation that handles the occlusion geometry for the entire floor. Set differences may not be easy to compute in general, but in the restricted case where the receiver is a line and the objects line segments, the geometry turns out to be very clean.

2.3 Polygonal Objects

Suppose now E and O are one-dimensional polygons (line segments) and the receiver is the line $y = 0$ in Flatland. In this case, the parameterization of $\mathbf{r}(t)$ is simply $\mathbf{r}(t) = (t, 0)$ and we simply use x for the position on the receiver. The transformation T becomes merely a translation by x and the subsequent projection $E \rightarrow E^*$ can be computed directly: if E has vertices (x'_i, y'_i) , $i = 1, 2$ then

$$p_i^* = \frac{x'_i - x}{y'_i}, \quad i = 1, 2.$$

Equation (2) can now be evaluated in closed form. Furthermore, each p_i^* as a function of x traces out a line, so the strip for M_E is an (unbounded) polygon.

To take occlusion by O into account, the strip M_O is constructed for O . The partially-occluded irradiance can be computed from

$$I(x) = S_{E^*}(x) - S_{E^* \cap O^*}(x),$$




Figure 4: (a) When endowed with the appropriate density function, the integral over each slice produces the irradiance (the smooth curve). (b) The occluded irradiance is the integral over the set difference of the emitter strip and the occluder strip. The removed portion, the intersection of the strips, is a parallelogram.




Figure 5: A non-parallel situation. (a) The strips are actually Flatland cones, but their intersection remains a four-sided quadrilateral. (b) The integral of the density function over cross-sections of the intersection (outlined in black) produces the subtractive shadow irradiance. One advantage of this geometry is that the density function is independent of x .

where

$$S_Q(x) = \int_{a(x)}^{b(x)} \omega(y) dy$$

in case Q is the line segment from $(a(x), 1)$ to $(b(x), 1)$.

The value of $S_{E \circlearrowleft O^*}(x)$ is therefore the integral of the irradiance density function $\omega(y)$ over the cut at x of the *intersection* of the strips $M_O \cap M_E$. This intersection has four vertices, one for each pair of edges of M_O and M_E . In the special case where O and E both happen to be parallel to the x -axis, M_O and M_E are both strips of constant width and the intersection $M_O \cap M_E$ is a parallelogram as shown in Figure 2. In the more general case, the intersection $M_O \cap M_E$ can be shown to be a quadrilateral, as in Figure 5. The quadrilateral is not self-intersecting as long as the line containing E do not intersect O , and vice-versa.

The formulation of irradiance as the integral over a cut of a two-dimensional object casts it as a weighted polyhedral spline. The constructions in this subsection show that both unoccluded and partially occluded irradiance have this formulation. What is the advantage? First of all, the formulation of irradiance in terms of a polyhedron encodes all the visibility information from the receiver in to a single, easily described two-dimensional object, as suggested in Figure 6. Second, and more importantly, polyhedral splines can be evaluated in closed form without actually finding the intersection of the cut [4].




Figure 6: (a) An occluder partitions space into regions of visibility. In this figure, the umbra does not reach the receiver. (b) The regions of visibility on the plane are stored in the intersection quadrilateral. The projection of the edges of the quadrilateral onto $y = 0$ give visibility information for each object vertex; the projected vertices are the points of discontinuity of the irradiance.

3 THE PERSPECTIVE PRISM

In this section we construct the perspective prism analogous to the one in flatland, which stores the apparent geometry of an arbitrary polygon in \mathbb{R}^3 at every point on the floor $z = 0$. The construction pertains to a scene element Q which is either an emitter or an occluder.

3.1 Definition by Cross-Section

Suppose Q is a planar polygon having vertices $\mathbf{v}_i = (x_i, y_i, z_i)$ for $i = 1, \dots, n$. Furthermore, assume Q is situated strictly above the floor, so that each $z_i > 0$. We define a subset M_Q of \mathbb{R}^4 in terms of its 2D cross-sections $\mathcal{P}_{(x,y)} \cap M_Q$, where for fixed (a, b)

$$\mathcal{P}_{(a,b)} = \{(a, b, z, w) : z, w \in \mathbb{R}\}$$

is an (a, b) -slice.

For each (x, y) the cross-section is given by the polygon $(\mathcal{V}_1, \dots, \mathcal{V}_n)$ including its interior, where

$$\mathcal{V}_i = (x, y, \frac{x_i - x}{z_i}, \frac{y_i - y}{z_i}), \quad i = 1, \dots, n. \quad (3)$$

We observe that the last two components of \mathcal{V}_i are the perspective transformation of (x_i, y_i) with the origin moved to (x, y) .

For later use we define Q^* as the polygon with vertices $((x_i - x)/z_i, (y_i - y)/z_i, 1)$ for $i = 1, \dots, n$. Thus $Q^* = Q^*(x, y)$ is a perspective transformation of Q including its interior to $z = 1$ with the origin moved to (x, y) .

3.2 The Geometric Structure of the Perspective Prism

Describing the cross-sections of M_Q is sufficient to properly define the prism as a point set, but the slices give only a local characterization, and, for our purpose, we need a quantitative understanding of the boundary of all of M_Q .

Suppose i is fixed. As (x, y) varies, the i th vertex $\mathcal{V}_i = \mathcal{V}_i(x, y)$ can be interpreted as a 2-surface in \mathbb{R}^4 that is in fact a 2-plane.

Each apparent edge of Q is a line segment that is a function of (x, y) , and can be parameterized by

$$\mathcal{E}_i(x, y, t) = (1 - t)\mathcal{V}_i(x, y) + t\mathcal{V}_{i+1}(x, y). \quad (4)$$




Figure 7: The product polyhedron $O \times E$ is formed by sweeping a copy of E (the red triangle) across O (the blue pentagon) or vice-versa. The boundary is thus formed by sweeping each object along each edge of the other, resulting in three-dimensional prisms. Occluder prism j shares a rectangular “side” with emitter prism i .

Equation (4) is a parameterized 3-surface in \mathbb{R}^4 with $x, y \in \mathbb{R}$ and $t \in [0, 1]$. That is, each apparent edge of Q “sweeps out” a 3-surface in \mathbb{R}^4 . Each slice is a polygon, and the boundary of the polygon consists exactly of its edges, so the boundary of M_Q is therefore the union of the sweep-surfaces of all the apparent edges. This shows that M_Q is actually a four-dimensional *manifold with boundary* embedded in \mathbb{R}^4 .

M_Q has even simpler structure if Q happens to be parallel to the floor. In this case, $z_i = h$ for each i , where h is the height of Q above the floor, and the vertex surface functions \mathcal{V}_i becomes

$$\begin{bmatrix} x \\ y \\ \frac{x_i - x}{z_i} \\ \frac{y_i - y}{z_i} \end{bmatrix} = \frac{1}{h} \begin{bmatrix} 0 \\ 0 \\ x_i \\ y_i \end{bmatrix} + \begin{bmatrix} x \\ y \\ -\frac{x}{h} \\ -\frac{y}{h} \end{bmatrix}$$

Thus M_Q can be viewed as Q rotated so that it is parallel to the zw -plane, scaled by $1/h$, and swept along a 2-plane in \mathbb{R}^4 . The edge surface functions, from (4)

$$\mathcal{E}_i(x, y, t) = \begin{bmatrix} x \\ y \\ \frac{x_i - x}{h} + t \frac{x_{i+1} - x_i}{h} \\ \frac{y_i - y}{h} + t \frac{y_{i+1} - y_i}{h} \end{bmatrix} \quad (5)$$

with the domain of t extended to all of \mathbb{R} , become proper hyperplanes (affine images of \mathbb{R}^3) in \mathbb{R}^4 . Therefore the perspective prism M_Q is a genuine, albeit unbounded, four-dimensional polyhedron.

4 THE OCCLUDER AND EMITTER PRISM

We now proceed to the situation where there is an occluding polygon in the scene. Suppose we have in \mathbb{R}^3 a uniformly emitting polygon $E = \langle e_1, \dots, e_m \rangle$ and an occluding polygon $O = \langle o_1, \dots, o_n \rangle$. We assume that E lies strictly above the plane of O , and also that both E and O lie strictly above the floor.

4.1 Algebra of Perspective Prisms

To analyze partially occluded emitters we need a 4D object M whose (x, y) cross-sections stores the geometry of the *visible* portion of the emitter. At each floor point (x, y) the *visible portion* of the apparent emitting polygon is simply $E^* - O^*$, where E^* and




Figure 8: The five emitter-prism facets of $O \times E$. Adjacent emitter prisms share a common emitter face.




Figure 9: The three occluder-prism facets of $O \times E$. Adjacent occluder prisms share a common occluder face.

O^* are E and O projectively transformed through (x, y) . From the definition of the perspective prism it follows immediately that

$$M = M_E - M_O.$$

Dealing with occlusion is now reduced to a tractable problem of performing booleans of the prisms.

4.2 The Structure of $O \times E$

If E and O are both parallel to the floor, then M_O and M_E are both four-dimensional “prism” polyhedra, and consequently their intersection is another polyhedron. It can be shown that the intersection of the two “prisms” is isomorphic to the product polyhedron $O \times E$, as in Flatland. In the more general case of non-parallel polygons the boundary structure is topologically equivalent to $O \times E$.

The product polyhedron $O \times E$ can be constructed by sweeping E , rotated to the zw -plane across all of O , or vice-versa. The boundary of $O \times E$ can be viewed as E swept along each edge of O , and vice-versa, as suggested in Figure 7. (Of course, the object is really four-dimensional, so these figures are only a representation.) Sweeping E along an edge of O produces a prism in the shape of E , a three-dimensional polyhedron embedded in \mathbb{R}^4 . Similarly, sweeping O along an edge of E produces an “occluder” prism. The boundary of $O \times E$ therefore consists of n emitter prisms, and m occluder prisms, each a three-dimensional polyhedron situated in \mathbb{R}^4 . These prisms are the $n + m$ *facets* of $O \times E$.

Each emitter prism has two emitter *faces*, which are two-dimensional polygons situated in \mathbb{R}^4 . Adjacent emitter prisms share a facet, and there is thus exactly one emitter face for each occluder vertex, and vice-versa. The rectangular *sides* of the prisms are also proper two-dimensional polygons situated in \mathbb{R}^4 . Occluder prism j shares exactly one side with emitter prism i , hence there are $m \times n$ sides. The n emitter faces, the m occluder faces, and the $m \times n$ sides together form the boundaries of the facets of $O \times E$.

5 IRRADIANCE

Our goal is to build a shadow function that produces exact irradiance everywhere. We review what the irradiance quantity is for a scene with polygonal occluders and emitters.

We recall that the irradiance from a uniformly emitting surface S , which is not self-occluding as viewed from a point (x_0, y_0) on a receiver, can be computed [1] from the surface integral

$$I(x_0, y_0) = K \int_S \frac{\cos \theta_0 \cos \theta}{\|r(x, y)\|^2} dS, \quad (6)$$

where $r(x, y) = (x - x_0, y - y_0)$ is the ray from (x_0, y_0) to a point (x, y) on S , θ_0 is the angle between the normal to the receiver at (x_0, y_0) and $r(x, y)$, while θ is the angle between $r(x, y)$ and the normal to S at (x, y) . The constant K is a physical parameter of S . For simplicity we hereafter assume that $K = 1$.

When the emitting surface is a polygon, the integral in (6) can be shown to be equal to

$$I(x_0, y_0) = \int_{E^*} \frac{1}{(1 + x^2 + y^2)^2} dx dy \quad (7)$$

where $E^* = E^*(x_0, y_0)$ is the perspective transformation of E onto the plane $z = 1$ with the origin moved to (x_0, y_0) . For each polygon Q^* in the plane $z = 1$ we define the function

$$S_{Q^*} = \int_{Q^*} \frac{1}{(1 + x^2 + y^2)^2} dx dy. \quad (8)$$

In an unoccluded scene with an emitter E the irradiance at a point (x_0, y_0) on the receiver is simply $I(x_0, y_0) = S_{E^*}(x_0, y_0)$. In a scene which in addition has an occluder O it follows from the algebra of perspective prisms that

$$I(x_0, y_0) = S_{E^* - O^*} = S_{E^*} - S_{E^* \cap O^*}. \quad (9)$$

This is another way of saying that the net irradiance at the point (x_0, y_0) on the receiver is obtained by subtracting the occluded irradiance from the full unoccluded quantity.

6 GENERALIZED POLYHEDRAL SPLINES

In this section we develop some mathematical tools that are needed to describe shadows. Suppose B is a bounded convex polyhedral body in \mathbb{R}^n whose boundary ∂B is a disjoint union of finitely many hyperplanes F_i , and suppose $\omega = \omega(\mathbf{y}) = \omega(y_1, \dots, y_n)$ is a function of n variables which is continuous and integrable over B . Given an integer $s \leq n$ let Q be the projection $\mathbb{R}^n \rightarrow \mathbb{R}^s$ given by $P(y_1, \dots, y_n) = (y_1, \dots, y_s)$. We define a generalized polyhedral spline $N_{B, \omega}(x)$ of s variables such that for all sufficiently smooth functions ϕ of s variables with local support

$$\int_{\mathbb{R}^s} N_{B, \omega}(\mathbf{x}) \phi(\mathbf{x}) d\mathbf{x} = \int_B \omega(\mathbf{y}) \phi(P\mathbf{y}) d\mathbf{y}. \quad (10)$$

Note that (10) defines $N_{B, \omega}$ as a distribution which we assume can be defined as a continuous function. If $n = s$ we then have

$$N_{B, \omega}(\mathbf{x}) = \chi_B(\mathbf{x}) \omega(\mathbf{x}), \quad (11)$$

where $\chi_B(\mathbf{x}) = 1$ if $\mathbf{x} \in B$ and 0 otherwise.

If $\omega = 1$ the function $N_{B, \omega}$ is the multivariate polyhedral spline [4, 32]. We will use more general densities ω for the definition of our shadow irradiance functions.

By writing the right-hand side of (10) as an iterated integral we obtain a common geometric interpretation of $N_{B, \omega}$

$$N_{B, \omega}(\mathbf{x}) = \int_{\mathbb{R}^{n-s}} \chi_B(\mathbf{x}, \mathbf{z}) \omega(\mathbf{x}, \mathbf{z}) d\mathbf{z}. \quad (12)$$

In this interpretation the value of the spline at a point $\mathbf{x} \in \mathbb{R}^s$ is obtained by integrating the density function ω over a $n - s$ -dimensional cross-section of B located at \mathbf{x} . This cross-section is the intersection of B and an \mathbf{x} -slice in \mathbb{R}^n .

To make the above formula tractable for evaluation we apply Stokes' theorem to (10) and obtain

$$N_{B, \omega}(\mathbf{x}) = \sum_i \text{sgn}(v_i) N_{P_{n-1}^n(F_i), \Omega \circ \Phi_i}(\mathbf{x}). \quad (13)$$

Here F_i is facet i of B , having outward normal \mathbf{n}_i , v_i is the n th component of \mathbf{n}_i , $\text{sgn}(t)$ is equal to 1, 0, -1 if $t > 0, t = 0, t < 0$, respectively, Ω is an anti-derivative of ω with respect to the last variable, $P_{n-1}^n(F_i)$ is the canonical projection of F_i to \mathbb{R}^{n-1} , and Φ_i is a function of $n - 1$ variables describing F_i in the form $\Phi_i(y_1, \dots, y_{n-1}) = (y_1, \dots, y_{n-1}, y_n)$, where y_n is an equation for the hyperplane F_i in terms of y_1, \dots, y_{n-1} . The sgn function handles a reversal of the orientation of facet F_i in the projection to the lower space.

We note that the functions on the right of (13) are splines of s variables defined using objects in \mathbb{R}^{n-1} . We can apply this formula recursively to reduce n . The recurrence stops when n reaches s , in which case the spline is given by (11).

The unweighted polyhedral spline recurrence ($\omega = 1$) developed by de Boor and Höllig [4], as well as the polynomially weighted recurrence of McCool [32] both are derived following an identity of Hakopian. Ours takes a different approach, based directly on Stokes' theorem. The main use for our formula is to derive the explicit shadow irradiance function in the next section.

6.1 Extension to Non-Polyhedral Objects

The recurrence (13) can be extended beyond proper polyhedra, to certain general manifolds having the same general structure as polyhedra, inasmuch as their boundaries consist of facet-like manifolds, each of one lower dimension. As we shall see, the general occluded polygonal irradiance problem leads to manifolds of this form. The primary difficulty of this extension is that the "facets" may not be expressible as explicit surfaces of their projections, and when they are, they are no longer the simple affine hyperplane functions as in the polyhedron case. Problems such as this can sometimes be circumvented by partitioning the boundary facets into disjoint pieces; however, in the object we are about to develop, it happens that we can effectively ignore this potential problem.

6.2 The Irradiance Function

In Section 5 we gave an explicit formula for the irradiance functions based on the geometry of prisms in \mathbb{R}^4 . In this section we have developed the appropriate spline theory to construct the irradiance functions. The remaining task is to combine the two approaches to show that the irradiance function is, in fact, a spline. First consider a scene with an unoccluded emitter. The irradiance at a point (x_0, y_0) on the receiver is given in (7), which we write in the form

$$I(x_0, y_0) = \int_{E^*} \omega(x_0, y_0, z, w) dz dw,$$

with

$$\omega(x_0, y_0, z, w) = \frac{1}{(1 + z^2 + w^2)^2}. \quad (14)$$




Figure 10: The shadow discontinuity lines consist of E projected through each point of O , and vice-versa.




Figure 11: The floor shadow structure for a non-parallel configuration. Here $n = 4$ and $m = 3$. (a) The three occluder images; each has a counter-clockwise orientation. (b) The four emitter images; each has a counter-clockwise orientation. (c) Two of the side images: $S_{1,1}$ has a negative (clockwise) orientation, $S_{3,1}$ is positively oriented.

Since E^* is essentially the cross-section of the prism M_E at (x_0, y_0) we obtain

$$I(x_0, y_0) = \int_{\mathbb{R}^2} \chi_{M_E}(x_0, y_0, z, w) \omega(x_0, y_0, z, w) dz dw.$$

Hence from (12) with $n = 4, s = 2$ we obtain

$$I(x_0, y_0) = N_{M_E} \omega(x_0, y_0), \quad (15)$$

and we have demonstrated that the irradiance function is a spline. This agrees with the fact that a shadow is a piecewise smooth function with jumps in derivatives across the discontinuity lines, and the nature of spline functions. The above derivation holds equally well in the more general situation including multiple elements. For example in the situation described in (9) we have

$$I(x_0, y_0) = N_{M_E - M_O} \omega(x_0, y_0). \quad (16)$$

7 THE SHADOW EQUATION

Using the recurrence relation (13) we can actually give a closed-form expression for the irradiance function corresponding to a scene with a single planar polygonal emitter and a planar polygonal occluder, in terms of the space coordinates of the vertices

$$\begin{aligned} O &= \langle (x_1, y_1, z_1), \dots, (x_n, y_n, z_n) \rangle \\ E &= \langle (x'_1, y'_1, z'_1), \dots, (x'_m, y'_m, z'_m) \rangle. \end{aligned}$$

We assume O and E are convex, and O lies strictly above the floor, and the plane of O lies strictly below E . We define the direction of the normal \mathbf{n}_E of E to be opposite the direction of emittance. (This is the opposite of the usual convention, but we need the emitter projections onto the plane to have counter-clockwise orientation.) The region on the plane in which the emitting side of E is visible is the set

$$R_E \equiv \{ (x, y) \in \mathbb{R}^2 : (x, y, 0) \cdot (\mathbf{n}_E - (x_1, y_1, z_1)) > 0 \} \quad (17)$$

The formula presented in this section is valid on the planar region $R_O \cap R_E$.

The shadow equation is expressed in terms of the characteristic functions of the projected 2D boundary faces of $O \times E$. To define these polygons, first define the *shadow nodes*

$$\mathbf{v}_{ij} \equiv \left(\frac{z'_j x_i - z_i x'_j}{z'_j - z_i}, \frac{z'_j y_i - z_i y'_j}{z'_j - z_i} \right). \quad (18)$$

Geometrically, \mathbf{v}_{ij} is the point on the floor where \mathbf{o}_i and \mathbf{e}_j appear to coincide. The shadow polygons consist of m occluder images O_j , n emitter images E_i , and mn quadrilaterals S_{ij} , as shown in Figures 10 and 11. In terms of the shadow nodes,

$$\begin{aligned} E_i &= \langle \mathbf{v}_{i,1}, \dots, \mathbf{v}_{i,m} \rangle \\ O_j &= \langle \mathbf{v}_{1,j}, \dots, \mathbf{v}_{n,j} \rangle \\ S_{ij} &= \langle \mathbf{v}_{i,j}, \mathbf{v}_{i+1,j}, \mathbf{v}_{i+1,j+1}, \mathbf{v}_{i,j+1} \rangle, \end{aligned}$$

The z and w coordinates of the vertices of the (x, y) -slices of the objects give the apparent geometry from (x, y) . Define (cf. Section 3.1.)

$$p_i(x, y) = (p_{i,1}, p_{i,2}) = \left(\frac{x_i - x}{z_i}, \frac{y_i - y}{z_i} \right) \quad (19)$$

$$p'_j(x, y) = (p'_{j,1}, p'_{j,2}) = \left(\frac{x'_j - x}{z'_j}, \frac{y'_j - y}{z'_j} \right). \quad (20)$$

It is convenient to introduce

$$\begin{aligned} e_i(x, y) &= (e_{i,1}, e_{i,2}) = p_{i+1}(x, y) - p_i(x, y) \\ e'_j(x, y) &= (e'_{j,1}, e'_{j,2}) = p'_{j+1}(x, y) - p'_j(x, y). \end{aligned}$$

For the evaluation of the spline, the function defining the i th facet must be written in the form $w_i(x, y, z) = m_i(x, y)z + b_i(x, y)$, where for O ,

$$\begin{aligned} m_i(x, y) &= \frac{e_{i,2}}{e_{i,1}} \\ b_i(x, y) &= p_{i,2} - \frac{e_{i,2}}{e_{i,1}} p_{i,1}, \end{aligned}$$

and for E ,

$$\begin{aligned} m'_j(x, y) &= \frac{e'_{j,2}}{e'_{j,1}} \\ b'_j(x, y) &= p_{j,2} - \frac{e'_{j,2}}{e'_{j,1}} p_{j,1}. \end{aligned}$$

The second iteration of (13) requires the equations for the boundary of the facets of M_O which are the 2-surfaces traced out by the vertices of O^* and E^* in M_O and M_E . These are essentially given by (3). There are also the extra boundary 2-surfaces traced by the intersection of edges i of O^* and j of E^* , which can be computed from

$$\psi_{ij}(x, y) = \frac{e_{i,1} e'_{j,2} p_{j,1} - e'_{j,1} (e_{i,2} p_{i,1} + e_{i,1} (p_{j,2} - p_{i,2}))}{e_{i,1} e'_{j,2} - e_{i,2} e'_{j,1}} \quad (21)$$

Finally, we need the twice-integrated density function

$$F(z, u, v) = \frac{v \arctan \left(\frac{uv + z + u^2 z}{\sqrt{1 + u^2 + v^2}} \right)}{2\sqrt{1 + u^2 + v^2}} + \frac{z \arctan \left(\frac{v + uz}{\sqrt{1 + z^2}} \right)}{2\sqrt{1 + z^2}} \quad (22)$$

The irradiance on the plane $z = 0$ is then

$$I(x, y) = N_{M_E} \omega(x, y) - N_{M_O \cap M_E} \omega(x, y)$$

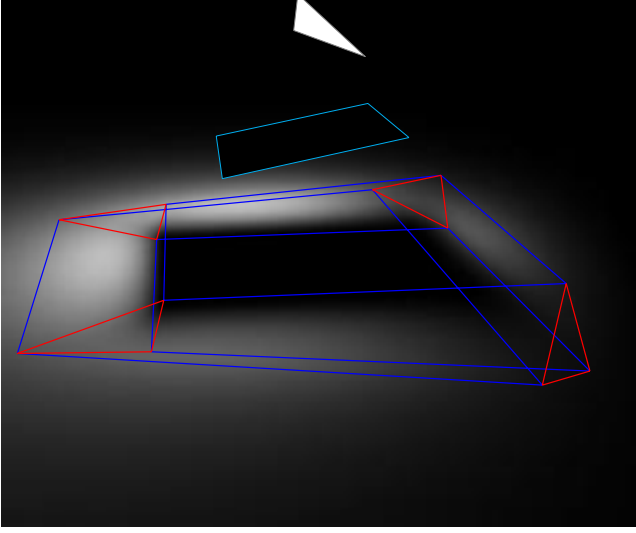


Figure 12: A non-parallel emitter/occluder configuration, showing the irradiance with discontinuity lines. Both E and O are outlined for clarity.

where ω is given by (14). Applying (13) twice it can be shown that

$$\begin{aligned}
 N_{M_O \cap M_E, \omega}(x, y) = & \\
 & \sum_{j=1}^m \chi_{\hat{O}_j}(x, y) [F(p'_{j,1}, m'_j, b'_j) - F(p'_{j,1}, m'_{j-1}, b'_{j-1})] \\
 + & \sum_{i=1}^n \chi_{\hat{E}_i}(x, y) [F(p_{i,1}, m_i, b_i) - F(p_{i,1}, m_{i-1}, b_{i-1})] \\
 + & \sum_{i=1}^n \sum_{j=1}^m o(S_{ij}) \chi_{\hat{S}_{ij}}(x, y) [F(\psi_{ij}, m_i, b_i) - F(\psi_{ij}, m'_j, b'_j)]
 \end{aligned} \tag{23}$$

Here \hat{Q} is the polygon Q excluding edges whose oriented angle with the x -axis is greater than zero, and less than π . The $o(Q)$ function is the orientation of the polygon Q : $o(Q)$ is 1 if the vertices are in counter-clockwise order, and -1 otherwise.

Inspection of the equation for $N_{M_O \cap M_E, \omega}(x, y)$ reveals the first summation is in fact the unoccluded irradiance, hence the irradiance can be expressed directly in the following shadow equation.

$$\begin{aligned}
 I(x, y) = & \\
 & \sum_{j=1}^m [1 - \chi_{\hat{O}_j}(x, y)] [F(p'_{j,1}, m'_j, b'_j) - F(p'_{j,1}, m'_{j-1}, b'_{j-1})] \\
 - & \sum_{i=1}^n \chi_{\hat{E}_i}(x, y) [F(p_{i,1}, m_i, b_i) - F(p_{i,1}, m_{i-1}, b_{i-1})] \\
 - & \sum_{i=1}^n \sum_{j=1}^m o(S_{ij}) \chi_{\hat{S}_{ij}}(x, y) [F(\psi_{ij}, m_i, b_i) - F(\psi_{ij}, m'_j, b'_j)].
 \end{aligned} \tag{24}$$

The shadow spline function $I(x, y)$ is therefore a collection of smooth functions supported on the shadow polygons O_j , E_i and $S_{i,j}$.

Images demonstrating the formula are shown in Figures 12, 13.




Figure 13: A typical difficult case, where the O (nearly) touches the floor. The direct shadow formula fails if the objects actually touch, but they may come arbitrarily close. Here the occluder has been lifted slightly, to reveal the C^0 discontinuity along the line of contact.

These images were rendered using ray tracing, with the irradiance on the surfaces calculated directly using the shadow equation.

The formula extends to multiple occluders as long as their shadows do not interact. The facets of a convex polyhedral occluder meet this, and the requirements of the shadow equation (if the normals point inward), so the formula can be extended to convex polyhedral emitters and occluders, as demonstrated in Figures 14 and 15.

8 FUTURE WORK AND CONCLUSION

The general formulation presented in this paper is rich in possibilities worthy of deeper investigation. As a first step, we hope to generalize the non-interfering occluder formula to take into account multiple objects and their interactions with each other. Beyond this, we are encouraged by the possibility of producing an analytic formula for objects more general than polygons, and receivers more general than portions of a plane. Furthermore, the ideas generalize to non-uniform emittance and arbitrary BRDFs.

8.1 Conclusion

In this paper we analyze the essential mathematical structure of shadows on a surface and show that the natural place to understand them is in a higher dimensional space. We show that the general problem of calculating the irradiance function for the shadow cast by a single polygonal occluder O from a polygonal emitter E in general position leads to a weighted (generalized) polyhedral spline function, and that the Cartesian product $O \times E$ provides the structure necessary to compute the shadow. To compute the irradiance for non-parallel configurations of E and O , ordinary polyhedral splines must be generalized. The resulting splines are built on objects which are deformations of polyhedra in higher dimensions, instead of the usual flat-faced structures.




Figure 14: The shadow of a regular dodecahedron cast by a triangular emitter. The irradiance is computed by subtracting the value of (23) for each face from the unoccluded irradiance; this works for convex polyhedra, because the edges do not interact under the emitter.

Using Stokes' theorem to reduce the Euclidean dimension of the integration problem, we arrive at a closed-form formula for evaluating the irradiance function, not just the partial visibility quantity, as an exact analytical expression. The wire frame structure of the 4D object by itself produces the shadow discontinuity lines when projected to the floor. That is, the projection method naturally specializes to produce the discontinuity graph described by Nishita and Nakamae [33]. When E , O , and the floor are all parallel and a constant weighting function is specified, the spline specializes to an $O \times E$ (prism) polyhedral spline that yields the convolution presented by Soler and Sillion in their shadow algorithm [38]. Although it is a common simplifying approximation to factor out visibility from the irradiance integral, there is no penalty in our method for doing it exactly.

Although exact shadows may not be necessary in many situations, they are important because their structure gives fundamental understanding. Knowing their exact behavior also provides a way to compute the error in approximate methods. And finally, the direct evaluation of the explicit formula is fast enough to compete with approximate algorithms.

A benefit of studying shadows in the higher dimensional context with the constructions provided is that the approach generalizes beyond the solution for a single emitter and occluder in arbitrary configuration. This approach also generalizes theoretically in a straightforward manner to more complicated situations involving curved and multiple emitter and occluder elements, at the expense of complicating the geometry. To carry forward this approach for curved and nonuniform elements we must again use generalized polyhedral splines.

Computing irradiance quantities for shadow functions is equivalent to computing the differential area to polygon form factor in the presence of an occluder. Although we do not take the space to develop the details, one can see that this approach may be used in radiosity. There is also a connection between the 4D emitter or




Figure 15: The shadow of a dodecahedron cast by three symmetrically placed emitters colored red, green, and blue. The total irradiance is the sum of the contributions from each emitter.

occluder prisms and some image based rendering algorithms.

Considering shadow irradiance functions as higher dimension generalized polyhedral splines has lead to several advances in understanding and computation. Although somewhat abstract, the approach described herein, which appears to be quite general and powerful, seems rather promising as a basis from which to do further work in this area and related problems.

Acknowledgments

This work was supported in part by DARPA (F33615-96-C-5621) and the NSF Science and Technology Center for Computer Graphics and Scientific Visualization (ASC-89-20219). The authors wish to thank John Hughes, Brian Smits, and Peter Shirley for reading earlier drafts of this paper, and the reviewers for their comments.

References

- [1] James Arvo. The irradiance Jacobian for Partially Occluded Polyhedral Sources. In Andrew Glassner, editor, *SIGGRAPH 94 Conference Proceedings*, Annual Conference Series, pages 343–350. ACM SIGGRAPH, ACM Press, July 1994. ISBN 0-89791-667-0.
- [2] James Arvo. *Analytic Methods for Simulated Light Transport*. PhD thesis, Yale University, 1995.
- [3] Daniel R. Baum, Holly E. Rushmeier, and James M. Winget. Improving Radiosity Solutions Through the Use of Analytically Determined Form-Factors. In Jeffrey Lane, editor, *Computer Graphics (SIGGRAPH 89 Conference Proceedings)*, volume 23, pages 325–334, July 1989. ISBN 0-89791-746-4.
- [4] C. de Boor and K. Höllig. Recurrence Relations for Multivariate B-Splines. *Proc. Amer. Math. Soc.*, pages 397–400, 1982.
- [5] C. de Boor, K. Höllig, and S. Riemenschneider. *Box Splines*. Springer-Verlag, 1993.
- [6] Arne Brøndsted. *An Introduction to Convex Polytopes*. Springer-Verlag, New York, 1983.

- [7] A. T. Campbell, III and Donald S. Fussell. Adaptive Mesh Generation for Global Diffuse Illumination. In Forest Baskett, editor, *Computer Graphics (SIGGRAPH 90 Conference Proceedings)*, volume 24, pages 155–164, August 1990. ISBN 0-89791-344-2.
- [8] Elaine Cohen, Tom Lyche, and Richard Riesenfeld. Discrete Box Splines and Refinement Algorithms. *Computer Aided Geometric Design*, 1(2):131–148, 1984.
- [9] Michael F. Cohen and Donald P. Greenberg. The Hemi-Cube: A Radiosity Solution for Complex Environments. In B. A. Barsky, editor, *Computer Graphics (SIGGRAPH 85 Conference Proceedings)*, volume 19, pages 31–40, August 1985.
- [10] Michael F. Cohen and John R. Wallace. *Radiosity and Realistic Image Synthesis*. Academic Press Professional, San Diego, CA, 1993.
- [11] Robert L. Cook, Thomas Porter, and Loren Carpenter. Distributed Ray Tracing. In Hank Christiansen, editor, *Computer Graphics (SIGGRAPH 84 Conference Proceedings)*, volume 18, pages 137–45, July 1984.
- [12] H. S. M. Coxeter. *Regular Complex Polytopes, Second Edition*. Cambridge University Press, New York, 1991.
- [13] Franklin C. Crow. Summed-area Tables for Texture Mapping. In Hank Christiansen, editor, *Computer Graphics (SIGGRAPH 84 Conference Proceedings)*, volume 18, pages 207–212, July 1984.
- [14] H. B. Curry and I. J. Schoenberg. On Pólya frequency functions IV: the fundamental spline functions and their limits. *J. Analyse Math.*, 17:71–107, 1966.
- [15] M. Dæhlen and T. Lyche. Box Splines and Applications. *Geometric Modelling, Methods and Applications*, pages 35–93, 1991.
- [16] W. Dahmen and C. A. Micchelli. Recent Progress in Multivariate Splines. In C. Chui, L. Schumaker, and J. Ward, editors, *Approximation Theory IV*, pages 27–121. Academic Press, New York, 1983.
- [17] George Drettakis. *Structured Sampling and Reconstruction of Illumination for Image Synthesis*. PhD thesis, University of Toronto, 1994.
- [18] George Drettakis and Eugene Fiume. A Fast Shadow Algorithm for Area Light Sources Using Backprojection. In Andrew Glassner, editor, *SIGGRAPH 94 Conference Proceedings*, Annual Conference Series, pages 223–230. ACM SIGGRAPH, ACM Press, July 1994. ISBN 0-89791-667-0.
- [19] P. Fong and Hans-Peter Seidel. Modeling with multivariate B-spline surfaces over arbitrary triangulations. In M. Silbermann and H. Tagare, editors, *Curves and Surfaces in Computer Vision and Graphics II*, pages 97–108. SPIE, 1992.
- [20] Philip Fong and Hans-Peter Seidel. An implementation of multivariate B-spline surfaces over arbitrary triangulations. In *Proceedings of Graphics Interface '92*, pages 1–10, May 1992.
- [21] Cindy M. Goral, Kenneth E. Torrance, Donald P. Greenberg, and Bennett Bataille. Modelling the Interaction of Light Between Diffuse Surfaces. In Hank Christiansen, editor, *Computer Graphics (SIGGRAPH 84 Conference Proceedings)*, volume 18, pages 212–22, July 1984.
- [22] Steven J. Gortler, Peter Schroder, Michael F. Cohen, and Pat Hanrahan. Wavelet Radiosity. In *Computer Graphics Proceedings, Annual Conference Series, 1993*, pages 221–230, 1993.
- [23] Branko Grünbaum. *Convex Polytopes*. John Wiley & Sons, 1967.
- [24] Pat Hanrahan, David Salzman, and Larry Uppperle. A Rapid Hierarchical Radiosity Algorithm. In Thomas W. Sederberg, editor, *Computer Graphics (SIGGRAPH 91 Conference Proceedings)*, volume 25, pages 197–206, July 1991. ISBN 0-89791-436-8.
- [25] Paul Heckbert. Discontinuity Meshing for Radiosity. *Third Eurographics Workshop on Rendering*, pages 203–226, May 1992.
- [26] Paul Heckbert. Radiosity in Flatland. *Computer Graphics Forum (Eurographics '92)*, 11(3):181–192, September 1992.
- [27] G. Heflin and G. Elber. Shadow Volume Generation from Free Form Surfaces. In *Communicating with Virtual Worlds, Proceedings of CGI'93 (Lausanne, Switzerland)*, pages 115–126. Springer-Verlag, June 1993.
- [28] J. R. Howell. *A Catalog of Radiation Configuration Factors*. McGraw Hill, 1982.
- [29] Marc Levoy and Pat Hanrahan. Light Field Rendering. In Holly Rushmeier, editor, *SIGGRAPH 96 Conference Proceedings*, Annual Conference Series, pages 31–42. ACM SIGGRAPH, Addison Wesley, August 1996. ISBN 0-89791-746-4.
- [30] Daniel Lischinski, Filippo Tampieri, and Donald P. Greenberg. Discontinuity Meshing for Accurate Radiosity. *IEEE Computer Graphics and Applications*, 12(6):25–39, November 1992.
- [31] Michael D. McCool. Analytic Antialiasing With Prism Splines. In Robert Cook, editor, *SIGGRAPH 95 Conference Proceedings*, Annual Conference Series, pages 429–436. ACM SIGGRAPH, Addison Wesley, August 1995. ISBN 0-89791-701-4.
- [32] Michael D. McCool. *Analytic Signal Processing for Computer Graphics using Multivariate Polyhedral Splines*. PhD thesis, University of Toronto, 1995.
- [33] Tomoyuki Nishita and Eihachiro Nakamae. Continuous Tone Representation of Three-Dimensional Objects Taking Account of Shadows and Interreflection. In B. A. Barsky, editor, *Computer Graphics (SIGGRAPH 85 Conference Proceedings)*, volume 19, pages 23–30, July 1985.
- [34] William T. Reeves, David H. Salesin, and Robert L. Cook. Rendering Antialiased Shadows with Depth Maps. In Maureen C. Stone, editor, *Computer Graphics (SIGGRAPH 87 Conference Proceedings)*, volume 21, pages 283–291, July 1987. ISBN 0-89791-227-6.
- [35] Peter Schröder and Pat Hanrahan. On the Form Factor Between Two Polygons. In James T. Kajiya, editor, *SIGGRAPH 93 Conference Proceedings*, Annual Conference Series, pages 163–164, 1993. ISBN 0-89791-601-8.
- [36] Peter Shirley, Chang Yaw Wang, and Kurt Zimmerman. Monte Carlo Techniques for Direct Lighting Calculations. *ACM Transactions on Graphics*, 15(1):1–36, January 1996. ISSN 0730-0301.
- [37] François Sillion and Claude Puech. *Radiosity and Global Illumination*. Morgan Kaufmann, San Francisco, 1994.
- [38] Cyril Soler and François X. Sillion. Fast Calculation of Soft Shadow Textures Using Convolution. In Michael Cohen, editor, *SIGGRAPH 98 Conference Proceedings*, Annual Conference Series, pages 321–332. ACM SIGGRAPH, Addison Wesley, July 1998. ISBN 0-89791-999-8.
- [39] Seth J. Teller. Computing the Antipenumbra of an Area Light Source. UCB/CSD 91 6, Computer Science Division, University of California, Berkeley, 1991.
- [40] John R. Wallace, Kells A. Elmquist, and Eric A. Haines. A Ray Tracing Algorithm for Progressive Radiosity. In Jeffrey Lane, editor, *Computer Graphics (SIGGRAPH 89 Conference Proceedings)*, volume 23, pages 315–324, July 1989.
- [41] Gregory J. Ward and Paul Heckbert. Irradiance Gradients. *Third Eurographics Workshop on Rendering*, pages 85–98, May 1992.
- [42] Turner Whitted. An Improved Illumination Model for Shaded Display. *Communications of the ACM*, 23(6):343–349, July 1980.
- [43] Lance Williams. Casting Curved Shadows on Curved Surfaces. In *Computer Graphics (SIGGRAPH 78 Conference Proceedings)*, volume 12, pages 270–274, Aug 1978.
- [44] Andrew Woo, Pierre Poulin, and Alain Fournier. A Survey of Shadow Algorithms. *IEEE Computer Graphics and Applications*, 10(6):13–32, November 1990.