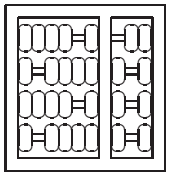

INSTITUTO DE COMPUTAÇÃO
UNIVERSIDADE ESTADUAL DE CAMPINAS

Comparison of Finite Element Bases for Global Illumination in Image Synthesis

Anamaria Gomide *Danillo R. Pereira*
Jorge Stolfi

Technical Report - IC-09-29 - Relatório Técnico

September - 2009 - Setembro

The contents of this report are the sole responsibility of the authors.
O conteúdo do presente relatório é de única responsabilidade dos autores.

Comparison of Finite Element Bases for Global Illumination in Image Synthesis

Anamaria Gomide, Danillo Roberto Pereira, Jorge Stolfi
Institute of Computing – University of Campinas (UNICAMP)
13083-970 , Campinas, SP, Brazil
anamaria@ic.unicamp.br,dpereira@ic.unicamp.br,stolfi@ic.unicamp.br *

Abstract

Finite element bases defined by sampling points were used by J. Lehtinen in 2008 for the efficient computation of global illumination in virtual scenes. The bases provide smooth approximations for the radiosity and spontaneous emission functions, leading to a discrete version of Kajiya's rendering equation. Unlike methods that are based on surface subdivision, Lehtinen's method can cope with arbitrarily complex geometries. In this paper we present an experimental validation of Lehtinen's meshless method by comparing its results with an independent numerical solution of the rendering equation on a simple three-dimensional scene. We also compare Lehtinen's special finite-element basis with two other similar bases that are often used for meshless data interpolation, namely a radial basis with a Gaussian mother function, and Shepard's inverse-square-distance weighted interpolation. The results confirm the superiority of Lehtinen's basis and clarify why the other two bases provide inferior-looking results.

1 Introduction

Realistic rendering usually requires modeling the indirect illumination, due to light that interacts two or more times with the scene's surface [1, 2]. For most scenes, the total light flow (including direct and indirect lighting) is adequately described by the *rendering equation* proposed by Jim Kajiya in 1986 [3].

Radiosity [2] is a general method for realistic rendering that uses finite element modeling to solve the rendering equation with Lambertian scenes. In this formulation, the surface of the scene is divided into a large number of *surface elements*. The light flow in the scene is found by solving a large system of linear equations $(I - R)\lambda = \varepsilon$, where the vector ε gives the spontaneous light emission and λ gives the total emission (spontaneous plus scattered) of each element.

Traditionally, the surface elements were the cells of a polygonal mesh approximating the scene's surface. A major source of difficulty in this approach is the complexity and variety of scene models, which called for rather complicated meshing algorithms. For one thing, computational cost often mandated the use of a mesh whose cells are much larger than many scene objects. Another source

*Institute of Computing – University of Campinas (UNICAMP), 13081-970. Research developed with financial support from CNPq.

of difficulty was the need to smooth out the inherent discontinuity of the radiosity between adjacent cells.

In 2007, Jaako Lehtinen proposed an alternative approach, where the surface patches are replaced by "fuzzy" finite elements, defined by a collection of sampling points on the surface [4, 5]. Lehtinen's *meshless radiosity* approach does not require an approximating mesh, but only the ability to find a point of the surface along a given ray. Therefore, it can cope with arbitrarily complex geometries, and can be used for almost any scene that can be rendered by ray tracing. Moreover, the representation is inherently smooth and provides fairly good results even with relatively coarse approximations.

Lehtinen's model for the radiosity function is a modified *radial basis* approximation [6], using a Gaussian-like kernel multiplied by a term that depends on the surface normal, and then adjusted to have the partition-of-unit property. Given the many approximations that are embedded in the method, its quantitative accuracy is not easy to analyze. Furthermore, there seems to be no published comparison of Lehtinen's basis with other finite-element bases that have often been used for meshless data interpolation.

In this paper we provide an experimental validation of Lehtinen's method by comparing its result on a simple scene with an independent numerical solution of the rendering equation. We also compare Lehtinen's basis with two other scattered-data interpolation methods, namely a radial basis [6] with Gaussian mother function, and Shepard's inverse-distance-squared interpolation formula [7]. For these two, we use a normal-sensitive distance function that captures Lehtinen's directional factor in a more systematic manner. The results validate Lehtinen's approach and provide insight on what qualities of a finite-element basis are most important for radiosity computations.

2 The rendering equation

Kajiya's rendering equation can be written as

$$L = E + \mathcal{R}L \tag{1}$$

where:

- $L(x, u)$ is the (unknown) *total radiance function*, the total light power emitted or scattered by the scene near the surface point x along directions near the unit vector u;
- $E(x, u)$ is the *spontaneous emittance function*, the light power emitted by the scene near x and along u;
- $\mathcal{R}$ is the *light transfer operator*, that expresses how light is scattered by the scene.

Informally, the *radiance* $L(x, u)$ is the apparent intensity of the light emitted by the surface near x, as seen by an observer in the direction u. It is the value that should be encoded in the corresponding pixel of a synthetic image of the scene, rendered assuming that the observer is in the direction u from x. The emittance $E(x, u)$ is the part of $L(x, u)$ that is due to light generated, rather than scattered, by the surface at x; it is therefore nonzero only on light sources that are part of the scene,

such as lamps; or parts that scatter light from sources external to the scene, such as a sunlit floor or wall.

The emittance and radiance functions usually depend on the wavelength band (color channel). However, in most applications one can render each channel independently. Therefore, in this paper we consider the rendering of single color channel, that is, with essentially monochromatic light.

The transfer operator $\mathcal{R}$ models the transport of light between the points of the scene and how that light interacts with the scene's objects. Its effect on an arbitrary function $Z(x,u)$ is

$$(\mathcal{R}Z)(x,u) = \int_{\mathbb{S}^2} \rho(x,v,u) Z(x{\uparrow}v, -v) H(x,v)\, dv \tag{2}$$

(see figure 1), where:

- $\mathbb{S}^2$ is the set of all directions (i.e. the unit sphere);
- $x{\uparrow}v$ is the first point of the scene's surface along the ray that leaves x in the direction v.
- $\rho(x,v,u)$ is the *bi-directional radiance distribution function* (BRDF), that gives the fraction of the incident light at the point x, coming from the direction v, that is re-emitted along directions near u;
- $H(x,v)$ is the *light spread factor* for the point x and direction v, that depends on the angle between the normal $n(x)$ of the surface at x and the light source direction v.

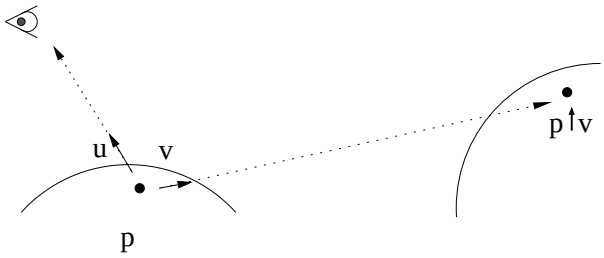

Figure 1: Parameters of the rendering equation, integrated over the direction v.

The factor $H(x,v)$ accounts for the fact that light coming from direction v gets spread over a larger or smaller portion of the scene surface at point p, depending on the angle between v and the surface normal $n(x)$ at x:

$$H(x,v) = n(x) \cdot v \tag{3}$$

Note that the visibility between the points of the surface is implicitly taken into account by the $\uparrow$ operation. Note also that this model allows translucent surfaces, if they are visible from both sides. By integrating over the scene's surface, instead over all directions, we get an alternative formulation of the light transfer operator $\mathcal{R}$:

$$\begin{aligned}(\mathcal{R}Z)(x,u) = \\ \int_{\mathcal{C}} \rho(x, x{\rightarrow}y, u) Z(y, x{\rightarrow}y) V(x,y) G(x,y)\, dy\end{aligned} \tag{4}$$

where

- $V(x, y)$ is the *visibility factor*, defined as 1 if light scattered or emitted at point x can illuminate point y, and 0 if that light is blocked by some other part of the scene strictly between x and y;
- $G(x, y)$ is the *geometric factor*, defined as $G(x, y) = H(x, x \to y)K(x, y)$ where H is the light spread factor (3), that accounts for the local inclination of the light reaching x from the direction of y, and K is the *apparent size factor*

$$K(x, y) = \frac{1}{4\pi} \frac{n(y) \cdot (y \to x)}{|x - y|^2} \tag{5}$$

  accounts for the apparent size (solid angle) of the area element dy as seen from x.

For an opaque Lambertian diffuse surface, the function $\rho(x, v, u)$ is $2\beta(x)$ if v and u are pointing out the object, and 0 otherwise; where $\beta(x)$ is the *scattering coefficient* or *albedo* ("intrinsic color") of the surface at the point x. See figure 2.

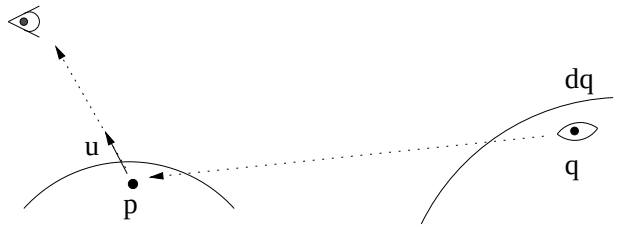

Figure 2: Parameters of the rendering equation, integrated over the surface point q.

The rendering equation, as described above, is still only an approximation to reality. It fails account for several physical phenomena, like diffraction, interference, polarization, and fluorescence. Fortunately, these phenomena are of little importance in the modeling of everyday ambients and objects.

2.1 Solving the rendering equation

In image synthesis, all elements of this equation are known except the radiance function L. Formally, the solution of the rendering equation is $L = (\mathcal{I} - \mathcal{R})^{-1}E$, where $\mathcal{I}$ is the identity operator and $^{-1}$ denotes operator inversion. In favorable circumstances, the *rendering operator* $(\mathcal{I} - \mathcal{R})^{-1}$ can be computed by Neumann's formula

$$(\mathcal{I} - \mathcal{R})^{-1} = (\mathcal{I} + \mathcal{R} + \mathcal{R}^2 + \mathcal{R}^3 + \cdots) \tag{6}$$

Each term $\mathcal{R}^k$ accounts for light that interacted k times with the surface of the scene before being observed.

3 Finite element radiosity

We use the term *site* to mean a pair $p = (\dot{p}, \vec{p})$ of a point $\dot{p}$ and a unit vector $\vec{p}$. We will denote by Γ the set of all sites, that is, $\mathbb{R}^3 \times \mathbb{S}^2$.

From now on we consider the emittance and radiance functions as being functions of pairs site-direction (p, v) instead of point-direction (x, v). We are only interested in the values of $E(p, v)$ and $L(p, v)$ for the subset $\mathcal{S}$ of all sites p that belong to the scene, namely where $\dot{p}$ is a point on the scene's surface and $\vec{p}$ is the corresponding unit normal vector. However, this formulation makes it possible to use a representation for the functions that is independent of the geometry of the surface.

For Lambertian radiosity, the emittance and radiance functions are independent of the direction v, so we can consider them as functions of the site p alone.

A *finite element* is a function ϕ defined on the scene's surface sites, such that $\phi(p)$ is nonzero only for a relatively small and compact set of sites (the *support* of ϕ). A *finite element basis* for radiosity is a collection $\phi = \{\phi_1, \phi_2, ...\phi_n\}$ of finite elements. A real-valued function of the surface sites, such as L or E, can be approximated by a linear combination $\mathcal{B}$ of basis elements

$$\mathcal{B}(p) = \sum_{i=1}^{n} \beta_i \phi_i(p) \tag{7}$$

where $\beta_1, \beta_2, \ldots \beta_n$ are real coefficients. For this purpose, the supports of the basis elements must cover the whole surface of the scene, and the functions ϕ_i must be linearly independent.

4 General properties of finite element bases

4.1 Interpolating bases

An *interpolation basis* is a function basis ϕ_1, ϕ_2,....ϕ_n with the property that $\psi(p_j)$ is 1 if $i = j$, and 0 otherwise. With such a basis, if f is any combination $\sum_i c_i \phi_i$, the value of f at each p_i is just c_i.

4.2 Partition-of-unity bases

We say that a basis $\phi_1, \phi_2, ..., \phi_n$ is a *partition of unity* if and only if

$$\phi_i(x) \geq 0 \tag{8}$$

for all i and all x in the domain Γ, and

$$\sum_{i=1}^{n} \phi_i(x) = 1 \tag{9}$$

for all $x \in \Gamma$. Such a basis has the *smoothing* property, namely

$$c_{\min} \geq \sum_{i=1}^{n} c_i \phi_i(x) \geq c_{\max} \tag{10}$$

where c_{min}, c_{max} are the minimum and maximum among the coefficients $c_1, c_2, \ldots, c_n$.

From any basis $\phi_1, \phi_2, ..., \phi_n$ with non-negative elements, one can define a basis with the partition-of-unity property $\tilde{\phi_1}, \tilde{\phi_2}, ...\tilde{\phi_n}$ by the *normalization* formula

$$\tilde{\phi_i}(x) = \frac{\phi_i(x)}{\sum_{j=1}^{n} \phi_j(x)} \qquad \text{for all } x \text{ in } \Gamma. \tag{11}$$

5 Generalized radial bases

The bases we use in this paper are *generalized radial bases*. A basis of this kind is defined by the following parameters:

- a list of sites $P = (p_1, p_2, ...p_n)$ on the scene's surface, the *element centroids*;
- a list of reals $(\alpha_1, \alpha_2, ...\alpha_n)$, the *nominal radii* of the elements;
- a *mother function* Φ from $\mathbb{R}$ to $\mathbb{R}$;
- a *distance metric* $\|\cdot, \cdot\|$ between sites;
- the element *scaling and placement* formula; and
- a *normalization* method applied to the basis elements.

In Lehtinen's method, the centroids are chosen randomly on the surface of the scene, as in figure 3; and each nominal radius α_i is such that there is a fixed number m of centroids $p_j \in P$ with $\|p_i, p_j\| < \alpha_i$. In this paper, we take $m = 10$, as used by Lehtinen.

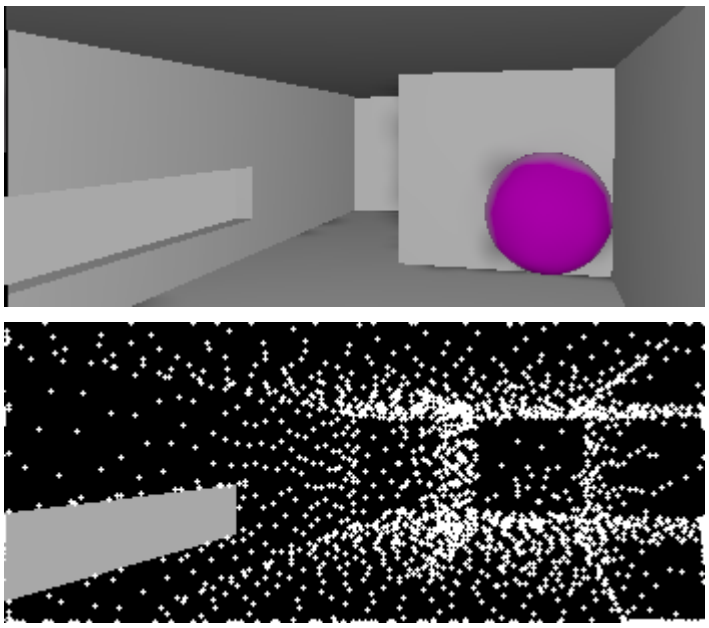

Figure 3: A simple scene (top) and a set of element centroids $p_1, \ldots, p_n$ randomly chosen on its surface (bottom). Note that invisible surfaces are sampled too.

The scaling and placement rules define a *raw element* ψ_i for each centroid, whose value $\psi_i(p)$ depends on the distance $\|p, p_i\|$, the radius α_i, and the mother function Φ. The raw elements are usually defined so that $\psi_i(p)$ is maximum (or nearly so) when $p = p_i$, and is zero (or nearly so) if $\|p, p_i\| > \alpha_i$. Finally, the normalization formula (11) may or may not be used to produce a partition-of-unity basis, yielding the final basis elements ϕ_i.

5.1 Mother functions

A commonly chosen mother function is the Gaussian bell

$$G(r) = e^{-r^2/2} \tag{12}$$

which, for the purposes of image synthesis, can be assumed to be zero for $r > 4$. For efficiency reasons, however, it is preferable to use a polynomial spline approximation

$$K(r) = \begin{cases} 2r^3 - 3r^2 + 1 & \text{if } r \leq 1 \\ 0 & \text{if } r \geq 1 \end{cases} \tag{13}$$

Another important alternative is the *Shepard's quadratic mother function*

$$S(r) = \frac{1}{r^2} \tag{14}$$

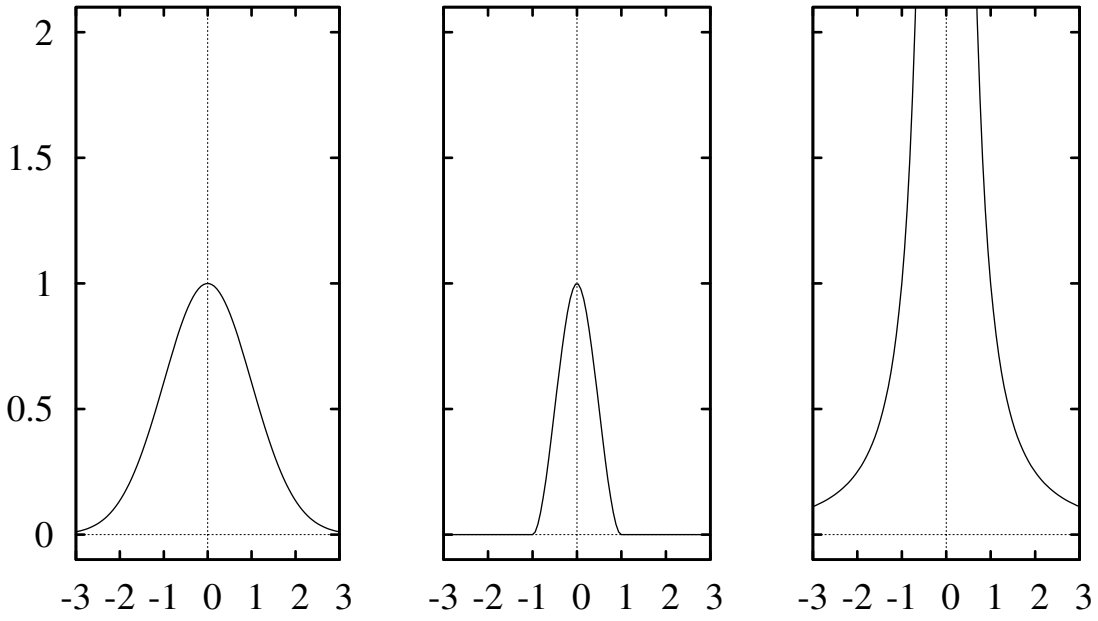

Figure 4: Three mother functions: Gaussian bell (left), spline bell (middle), and quadratic Shepard (right).

5.2 Site distance function

In typical scenes, the radiance (apparent color) of most points depends strongly on the local orientation of the surface. For this reason, one should take the normals into account when interpolating the radiance, at a point p, so that centroids with the same orientation as p get more weight than

centroids that are closer to p but have different orientation. Therefore, when computing the distance between two sites p and q, we use the *normal sensitive site distance*

$$\|p, q\| = \frac{|\dot{p} - \dot{q}|}{\max\{0, \vec{p} \cdot \vec{q}\}} \tag{15}$$

where $\vec{p} \cdot \vec{q}$ is the scalar product of the two normals. (This not a true distance function (metric) for Γ, because it fails the triangle inequality; however, that property is not necessary for interpolation.) The difference between formula (15) and the plain Euclidean distance $|\dot{p} - \dot{q}|$ is illustrated in figure 5.

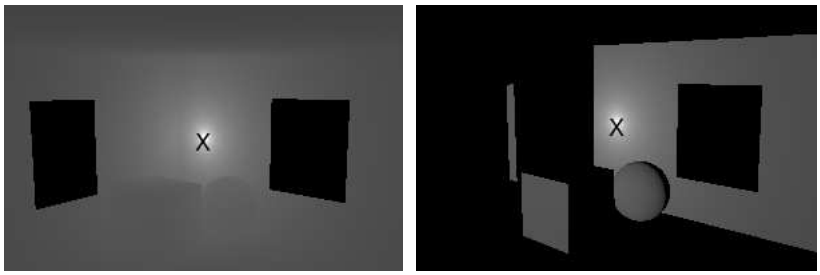

Figure 5: Visualization of the Euclidean distance (left) and normal-sensitive site distance (right) in a simple scene. The color used at each site p of the scene's surface is $1/(1 + d)$, where d is the distance from the site q on the back wall marked with 'X'.

6 Bases used in the tests

In our tests we used three generalized radial bases ϕ^{S}, ϕ^{G}, and ϕ^{L}. The basis ϕ^{S} uses the Shepard mother function S (14), applied to the absolute site distance, with partition-of-unit normalization:

$$\psi_i^{\mathrm{S}}(p) = S(\|p, p_i\|) = \frac{1}{\|p, p_i\|^2}; \quad \phi_i^{\mathrm{S}}(p) = \frac{\psi_i^{\mathrm{S}}(p)}{\sum_{j=1}^{n} \psi_j^{\mathrm{S}}(p)} \tag{16}$$

Note that $S(r) = 1/r^2$ is positive for all r and tends to $+\infty$ when r approaches 0. This property together with formula (16) ensures that Shepard's basis is always interpolating. See figure 6(top).

The basis ϕ^{G}, shown in figure **??**(middle), uses the Gaussian mother function (12), applied to the relative site distance from the centroid (site distance divided by the nominal radius), without partition-of-unit normalization:

$$\psi_i^{\mathrm{G}}(p) = G\left(\frac{\|p, p_i\|}{\alpha_i}\right); \quad \phi_i^{\mathrm{G}}(p) = \psi_i^{\mathrm{G}}(p) \tag{17}$$

The basis ϕ^{L} is the basis described by Lehtinen **??**. It uses the polynomial mother function K (13), but applied to the relative Euclidean distance, instead of our normal-sensitive site distance, with an external factor to account for the difference in the normals:

$$\psi_i^{\mathrm{L}}(p) = K\left(\frac{|\dot{p},\, \dot{p}_i|}{\dot{\alpha}_i}\right) \max\left\{0, \vec{p}\langle|\vec{p}_i\rangle\right\} \tag{18}$$

Here $\dot{\alpha}_i$ is computed like our radius α_i, but with Euclidean distances instead of site distance. The raw basis ψ^{L} was then normalized by formula 11 to yield a partition-of-unity basis ϕ^{L} See figure 6(bottom).

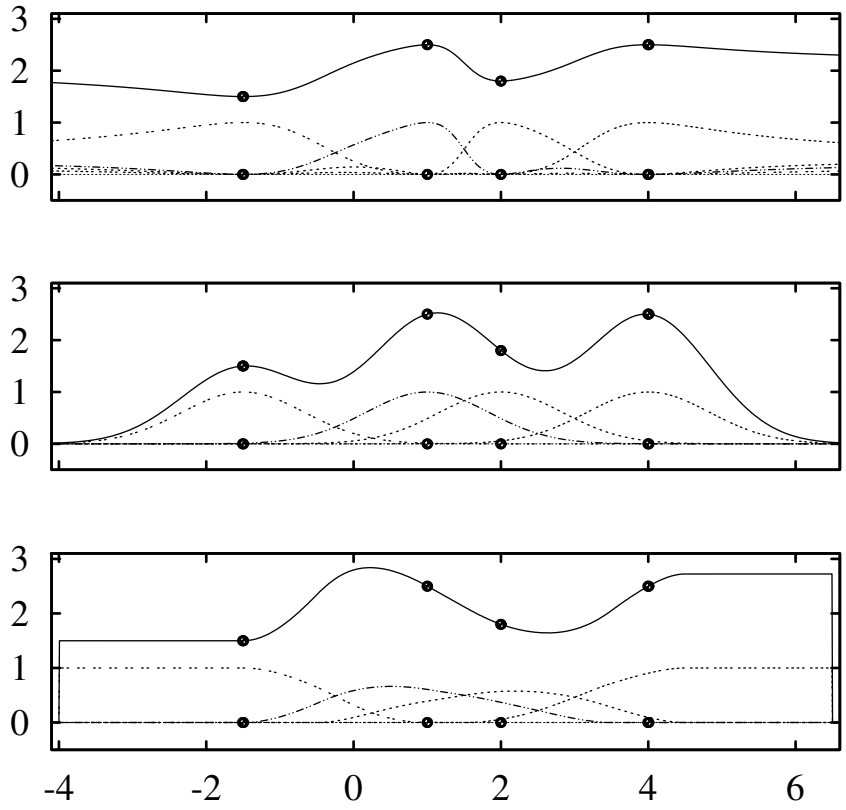

Figure 6: One-dimensional plots of the three test bases ϕ^{S}, ϕ^{G}, and ϕ^{L} (light lines) for four collinear sites on a flat surface, and the corresponding interpolation $f(x)$ to the four values shown (solid line).

7 Discretizing the rendering equation

When the functions L, E of the rendering equation (1) are represented in terms of a finite bases $\phi_1, \phi_2, ..., \phi_n$, the transfer operator $\mathcal{R}$ is replaced by an $n \times n$ *radiance transfer matrix* R and the equation becomes a linear equation system

$$\lambda = \varepsilon + R\lambda \tag{19}$$

where $\lambda = (\lambda_1, \lambda_2, \ldots, \lambda_n)^\top$ is the column coefficient vector of L in the chosen basis, and $\varepsilon = (\varepsilon_1, \varepsilon_2, \ldots, \varepsilon_n)^\top$ is the coefficient vector of E.

Each element R_{jk} of R represents the fraction of photons radiated (emitted or scattered) by element ϕ_k that are subsequently scattered by element ϕ_j, without any intermediate scattering. That is, column k of R is the coefficient vector λ that describes the appearance of the scene when element ϕ_k is the only source of light in the scene, without considering multiply-scattered light. See figure 7.

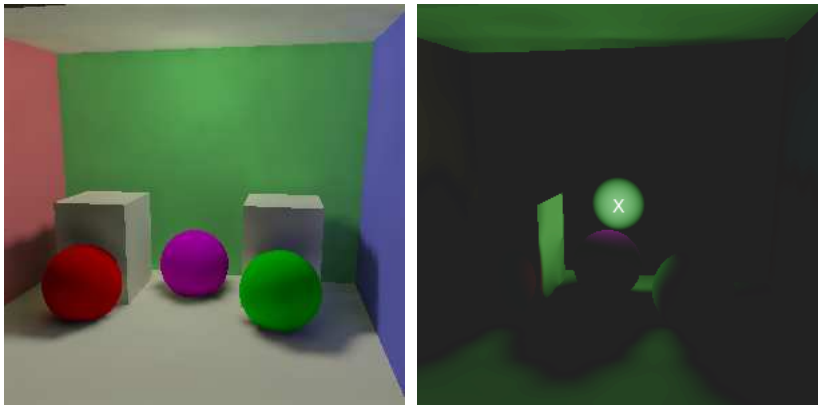

Figure 7: A simple scene (top), and the radiance L of its surface due to single-scattered photons emitted by the basis element ϕ_k whose centroid is marked with 'X' (bottom).

Lehtinen observed that one can estimate the matrix R by assuming a point light source of appropriate intensity w_j located at each site p_j, and computing the radiance $N_{i,j} = L(p_i)$ due to single-bounce photons from that source, as in plain ray-tracing. The intensity w_j of the point light can be approximated by the total radiance of the element ϕ_j, that is $w_j = \int \phi_j(p)dp$ where the integral is taken over the whole surface of the scene. If the centroids are sufficiently dense, we can assume that the scene surface near p_j is a plane with normal $\vec{p}_j$. For an un-normalized radial basis like ϕ^{G}, the integral is a fixed constant times α_j^2. For a partition-of-unity basis like ϕ^{L} and ϕ^{S}, the expected value of the integral is $1/\delta$ where δ is the local density of centroids per unit of area. If the radius α_j is chosen so that it contains t other centroids, then we can use the estimate $w_j = 1/\delta = \pi\alpha_j^2/t$.

The matrix R is not very sparse in general, and the inverse $(I - R)^{-1}$ is usually full. Therefore, the coefficients λ_i of the radiance function L are usually computed iteratively, by setting $\lambda \leftarrow (0, ..., 0)$ and then iterating $\lambda \leftarrow \varepsilon + R\lambda$ until convergence. Note that this iteration is equivalent to evaluating Neumann's formula (6).

Figure 8 shows a test scene rendered with plain ray-tracing and with meshless finite-element radiosity, using 10 iterations of formula (19).

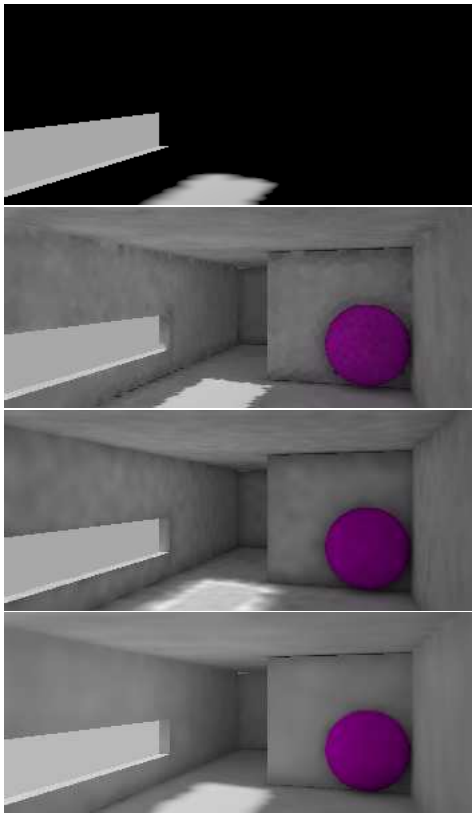

Figure 8: Images of a test scene. From top: the direct lighting component (represented in the ϕ^{L} basis), and the meshless radiosity results using the bases ϕ^{S}, ϕ^{G}, and ϕ^{L}.

8 Validation

In order to validate our implementation of Lehtinen's method and compare its accuracy with the three bases, we compared its results with an independent solution of the radiance equations for a specific scene. The latter consists of a sphere of radius r and a disk of radius R with the center of the ball on the disk's axis, at some distance h from its top surface. See figure 9(left). Both objects have Lambertian finish, with uniform albedo β. The primary illumination in this test case is due to a single point source with intensity μ, on the vertical axis, at infinite distance above the disk.

Let $p(\theta, \zeta)$ be the point on the sphere's surface at longitude θ and latitude ζ; and let $q(\varphi, u)$ be the point on the cylinder's top surface at distance u from the center and azimuth φ. Due the symmetry of the scene and lighting around the vertical axis, we can conclude that the radiance functions L and E are also symmetric (independent of the azimuths θ and φ).

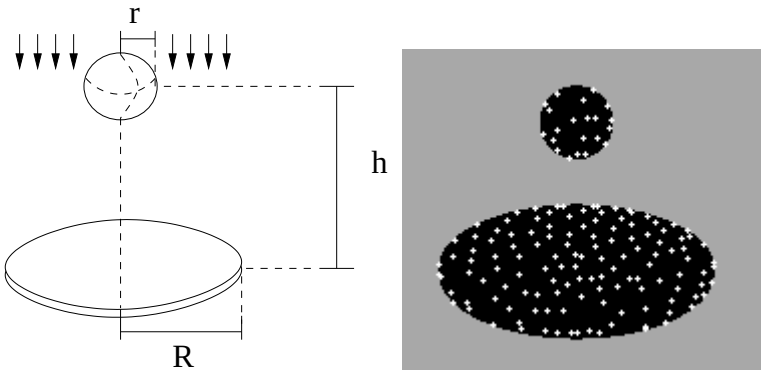

Figure 9: The reference scene (left) and the element centroids used in the meshless radiosity computation (right).

Therefore we denote by $B(\zeta)$ the total radiance (apparent color) of the sphere point $p(\theta, \zeta)$, and by $D(u)$ the radiance of the disk point $q(\varphi, u)$. It is convenient to consider the photons that have been scattered only once as being emitted at the scattering point, so that the external source can be ignored. We will denote this "emission" component of B and D as $B^*(\zeta)$ and $D^*(u)$, respectively:

$$\begin{aligned} B^*(\zeta) &= \beta\mu \max\{0\,,\,\sin\zeta\} \\ D^*(u) &= \begin{cases} \beta\,\mu & \text{if } r \le u \le R \\ 0 & \text{otherwise} \end{cases} \end{aligned} \tag{20}$$

From symmetry it also follows that the form factor $F(q(\varphi, u), p(\theta, \zeta))$ can be written $\hat{F}(u, \tau, \zeta) = F(q(0, u), p(\tau, \zeta))$ where $\tau = \theta - \varphi$. With these assumptions, the rendering equation can be rewritten as two coupled integral equations

$$\begin{aligned} B(\zeta) &= B^*(\zeta) + \beta \int_0^R D(u) \int_0^{2\pi} \hat{F}(u, \tau, \zeta)\, u\, d\tau\, du \\ D(u) &= D^*(u) + \beta \int_{-\pi/2}^{\pi/2} B(\zeta) \int_0^{2\pi} \hat{F}(u, \tau, \zeta)\, r^2\, \cos\zeta\, d\tau\, d\zeta \end{aligned} \tag{21}$$

Note that the visibility factor of $V(q(\varphi, u), p(\theta, \zeta))$ is 0 only if the ball normal at the point $p(\theta, \zeta)$ makes an obtuse angle with the direction $q(\varphi, u) \rightarrow p(\theta, \zeta)$; but in this case the form factor F is 0. Therefore, we do not need to include V in these formulas.

8.1 Discretization of the reference solution

In order to discretize the equations (21), we choose latitudes $\zeta_1, \ldots, \zeta_m$ in the interval $[-\pi/2\,,\,\pi/2]$, radii $u_1, \ldots, u_n$ in $[0\,,\,R]$, and azimuth differences $\tau_1, \ldots, \tau_k$ in $[0\,,\,2\pi]$, all equally spaced, and introduce the unknowns $b_i = B(\zeta_i)$, $d_j = D(u_j)$, and the known parameters $b_i^* = B^*(\zeta_i)$, $d_j^* = D^*(u_j)$, $F_{isj} = F(p(0, \zeta_i), q(\tau_s, u_j))$ and $G_{jsi} = F(q(0, u_j), p(\tau_s, \zeta_i))$. Then the integrals (21) and (**??**) can be approximated by sums:

$$b_i \;=\; b_i^* + \beta \frac{R}{n} \sum_{j=1}^{n} d_j \frac{2\pi}{k} \sum_{s=1}^{k} F_{isj}\, u_j \tag{22}$$

$$d_j \;=\; d_j^* + \beta\frac{\pi}{m}\sum_{i=1}^{m} b_i \frac{2\pi}{k}\sum_{s=1}^{k} G_{jsi}\, r^2\, \cos\zeta_i \tag{23}$$

We can write the equations (22) and (23) in the matrix form $\lambda = \varepsilon + R\,\lambda$, where $\lambda = (b_1, b_2, \ldots, b_m, d_1, d_2, \ldots, d_n)^\top$, $\varepsilon = (b_1^*, b_2^*, \ldots, b_m^*, d_1^*, d_2^*, \ldots, d_n^*)^\top$,

$$R = \begin{pmatrix} 0 & M \\ N & 0 \end{pmatrix} \tag{24}$$

and

$$N_{ji} = \beta\sum_{s=1}^{k} F_{isj} \qquad M_{ij} = \beta\sum_{s=1}^{l} G_{jsi} \tag{25}$$

Each element N_{ji} represents the influence of the radiance of band j of the disk on the radiance of each point of the ring i of the sphere. Similarly, M_{ij} represents the influence of ring i of the sphere at each point of ring j of the disk.

The parameters we used were $R = 40$, $r = 5$, $h = 20$, $\mu = 0.9$, $\beta = 0.9$, $n = m = k = 100$. We solved the system by iterating $\lambda \leftarrow \varepsilon + R\lambda$ (which converged after a few iterations). See figure 10

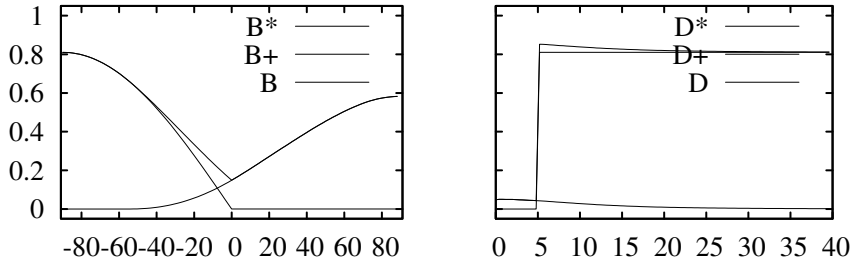

Figure 10: Reference solution, showing the radiance $B(\zeta)$ on the sphere as a function of latitude ζ (left), and radiance on the disk $D(u)$ as a function of the radial position u (right).

8.2 Comparison

The reference solution obtained as described in section 8.1 was compared with the output of the meshless radiosity algorithm described in section 7, using each of the three bases described in section 6, for the set of centroids shown in figure 9(right), chosen so that their minimum separation is 3. The results are shown in figures 11, 12, and 13. The thin line show the radiance along the meridian with latitude $\theta = 0$ of the sphere, and along the ray with azimuth $\varphi = 0$ of the disk. The thick lines are the radiances averaged over all latitudes θ and all azimuths φ, namely over each parallel of the sphere and each circle on the disk.

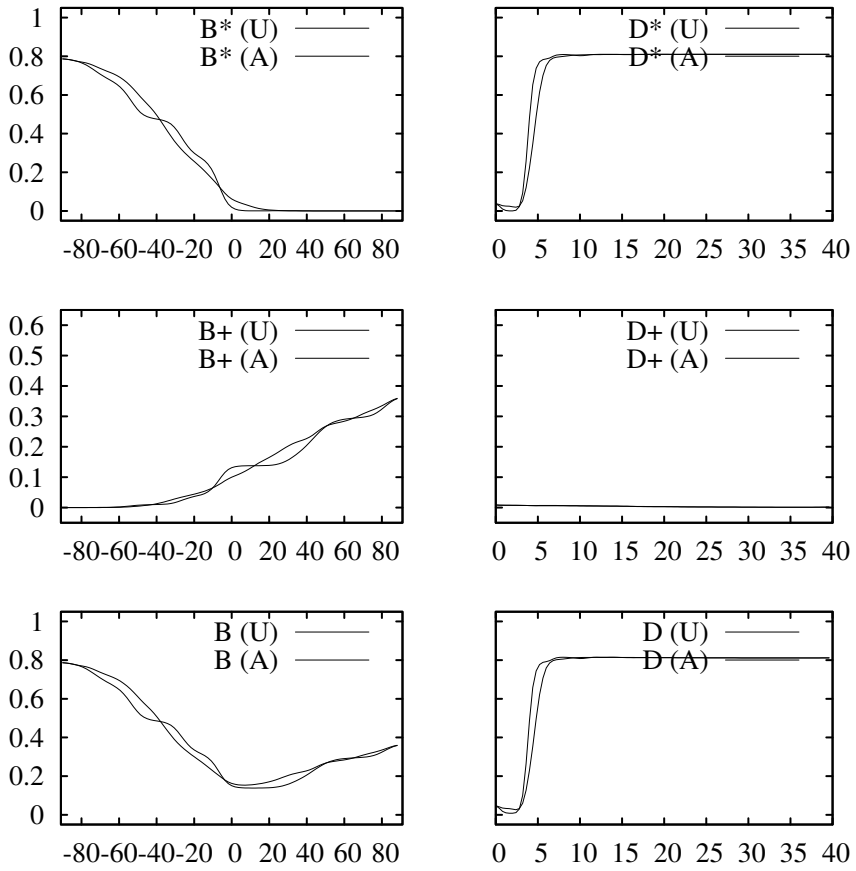

Figure 11: Meshless radiosity solution with the basis ϕ^{S}

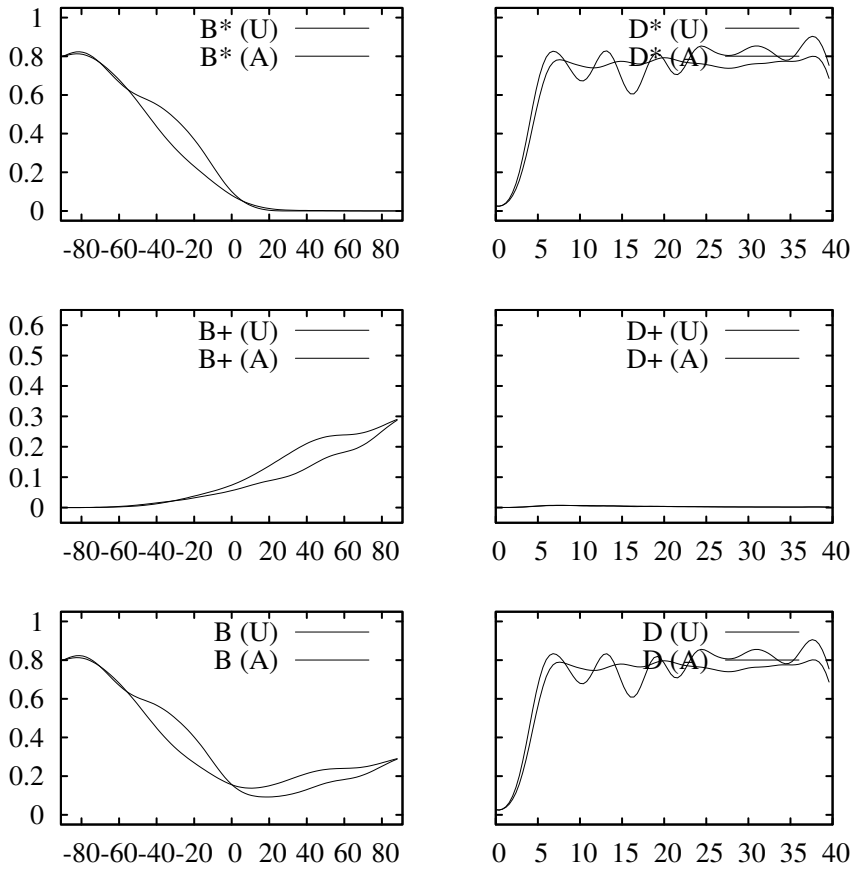

Figure 12: Meshless radiosity solution with the basis ϕ^{G}

Further tests (not shown here) imply that the radiance computed with any of the three bases tends to the same values as the number of elements increases.

9 Conclusions

Comparing the solutions obtained with the three bases with the reference solution, we conclude that Lehtinen's basis ϕ^{L} not only produce better-looking images (section 7) but also more accurate radiance values (section 8). We believe that the validation strategy proposed in section 8 can be used for the validation of other global illumination algorithms.

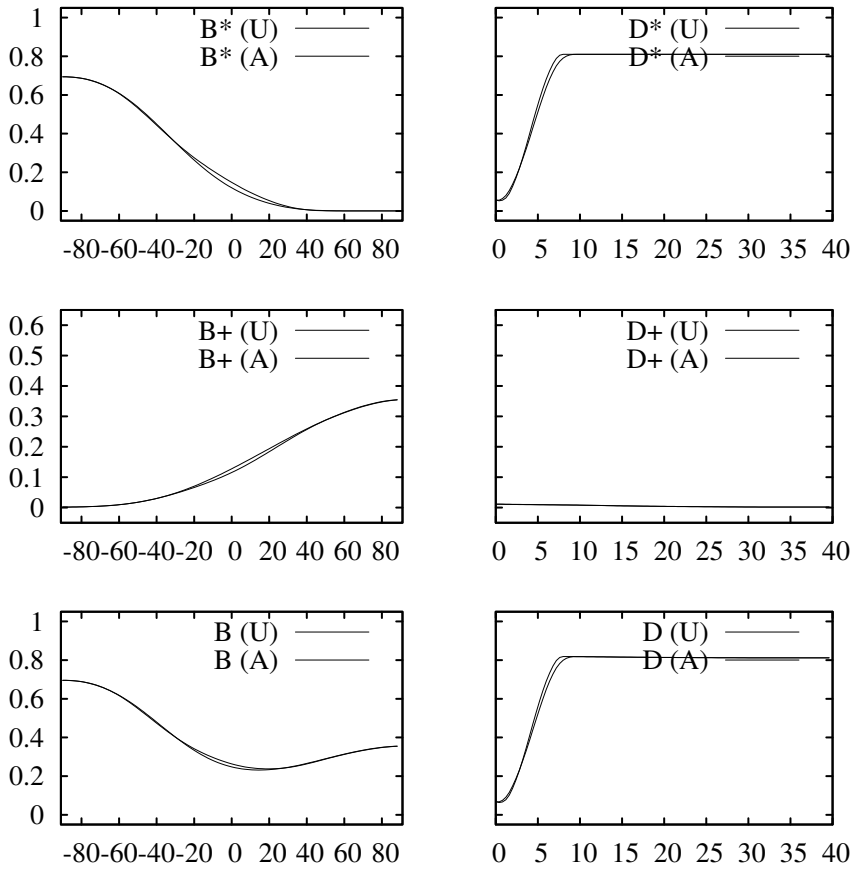

Figure 13: Meshless radiosity solution with the basis ϕ^{L}

Our analysis also shows that the inferior-looking results obtained by the other two bases are not due to systematic errors in the computed radiance values, but rather to the oscillations that they introduce in the approximation, which are magnified by the human visual system.

References

[1] T. Whitted, "An improved illumination model for shaded display," *Communications of the ACM*, vol. 23, no. 6, pp. 343–349, June 1980.

[2] M. F. Cohen and J. R. Wallace, *Radiosity and Realistic Image Synthesis*. Academic Press, 1993.

[3] J. T. Kajiya, "The rendering equation," in *Proc. SIGGRAPH 1986*, 1986, pp. 143–150.

[4] J. Lehtinen, M. Zwicker, J. Kontkanen, E. Turquin, F. Sillion, and T. Aila, "Meshless finite elements for hierarchical global illumination," Helsinki University of Technology, Tech. Rep. TML-B7, May 2007. [Online]. Available: http://artis.imag.fr/Publications/2007/LZKTSA07

[5] J. Lehtinen, M. Zwicker, E. Turquin, J. Kontkanen, F. Durand, F. Sillion, and T. Aila, "A meshless hierarchical representation for light transport," *ACM Trans. Graph.*, vol. 27, no. 3, 2008.

[6] M. S. Floater and A. Iske, "Multistep scattered data interpolation using compactly supported radial basis functions," *J. Comput. Appl. Math.*, vol. 73, no. 1-2, pp. 65–78, 1996.

[7] D. Shepard, "A two-dimensional interpolation function for irregularly-spaced data," in *Proc. 23rd ACM National Conference*, 1968, pp. 517–524.