

Wavelet Radiance[†]

Per H. Christensen Eric J. Stollnitz Tony D. DeRose David H. Salesin

University of Washington

Abstract

In this paper, we show how wavelet analysis can be used to provide an efficient solution method for global illumination with glossy and diffuse reflections. Wavelets are used to sparsely represent radiance distribution functions and the transport operator. In contrast to previous wavelet methods (for radiosity), our algorithm transports light directly among wavelets, and eliminates the pushing and pulling procedures.

The framework we describe supports curved surfaces and spatially-varying anisotropic BRDFs. We use importance to make the global illumination problem tractable for complex scenes, and we use a final gathering step to improve the visual quality of the solution.

1 Introduction

Radiosity algorithms assume that all reflection is ideally diffuse. This assumption, while making the computation of global illumination more tractable, ignores many important effects, such as glossy highlights and mirror reflections. Though more expensive, the simulation of directional reflection is essential for realistic image synthesis.

One promising approach to solving directional light transport is the finite element method, as pioneered by Immel *et al.* [15] and Shao *et al.* [20], and later refined by Sillion *et al.* [21]. Recently, Gortler *et al.* [13] and Schröder *et al.* [19] proposed an algorithm based on wavelets that focuses effort on the significant energy transfers, for the simpler case of radiosity. These works use the “non-standard” decomposition of the transport operator, and represent radiosity as a weighted sum of scaling functions. In addition, they require the use of “Push” and “Pull” procedures to distribute radiosity among levels of a hierarchy in each iteration.

Building on this work, we have developed a four-dimensional wavelet representation for spatially- and angularly-varying radiance distributions. However, in contrast to the approach taken by Gortler *et al.*, our algorithm uses the “standard” decomposition of the transport operator, and represents radiance in terms of wavelets rather than scaling functions. Our method does not require the pushing and pulling procedures.

We also incorporate importance-driven refinement, as described by Smits *et al.* [22] for radiosity, to avoid unnecessary work in computing viewpoint-dependent solutions of complex scenes. In addition, our implementation supports curved surfaces and anisotropic bidirectional reflectance distribution functions. The framework we describe naturally accommodates spatial variations, described by texture maps, in both emission and reflectance. Finally, to improve the visual quality of the image, a final gathering step is used [18].

[†]Per H. Christensen, Eric J. Stollnitz, David H. Salesin, and Tony D. DeRose. Wavelet radiance. In G. Sakas, P. Shirley, and S. Miller, editors, *Photorealistic Rendering Techniques* pages 295–309. Springer-Verlag, Berlin, 1995.

Reprinted from *Proceedings of the Fifth Eurographics Workshop on Rendering* (Darmstadt, Germany, June 1994), pages 287–302.

2 Finite Elements for Radiance

In this section, we briefly review the equation that governs light transport, and describe how the finite element method can be used to compute approximate solutions.

2.1 Radiance

Let x , y , and z be points in space. *Radiance* $L(y \rightarrow z)$ is defined as the power emanating from y , per unit solid angle in the direction towards z , per unit projected area perpendicular to that direction.

At equilibrium, radiance satisfies the following transport equation [10]:

$$L(y \rightarrow z) = L_e(y \rightarrow z) + \int_x f_r(x, y, z) G(x, y) L(x \rightarrow y) dx. \quad (1)$$

In this equation, $L_e(y \rightarrow z)$ is the *emitted radiance* from y in the direction towards z , and dx is an infinitesimal area around point x . The term $f_r(x, y, z)$ is the *bidirectional reflectance distribution function*, or BRDF, describing the ratio of reflected radiance (in the direction towards z) to the differential irradiance (from the direction of x) that produces it. Finally, the *geometric term* $G(x, y)$ is given by

$$G(x, y) \equiv V(x, y) \cdot \frac{\cos \theta_x \cos \theta_y}{\|x - y\|^2},$$

where $V(x, y)$ is a *visibility term* that is 1 or 0, depending on whether or not x and y are visible to one another, and θ_x and θ_y are the angles between the line segment xy and the respective normals of differential areas at x and y . The geometric term describes how radiance leaving a differential area at x in the direction towards y arrives as differential irradiance at y .

The transport equation (1) can be rewritten in operator form as

$$L = L_e + \mathcal{T}L. \quad (2)$$

Here, the *transport operator* \mathcal{T} is defined by

$$(\mathcal{T}L)(y \rightarrow z) \equiv \int_x f_r(x, y, z) G(x, y) L(x \rightarrow y) dx,$$

where $(\mathcal{T}L)(y \rightarrow z)$ denotes the result of \mathcal{T} operating on $L(x \rightarrow y)$ to produce a function whose argument is $(y \rightarrow z)$.

2.2 Discretization

Let $\mathbf{B}(x \rightarrow y) = (b_1(x \rightarrow y), b_2(x \rightarrow y), \dots)$ be a basis for the space of radiance distributions. The unknown radiance distribution can be projected onto the basis \mathbf{B} by writing L as a series expansion,

$$L(x \rightarrow y) = \sum_{i=1}^{\infty} \ell_i b_i(x \rightarrow y).$$

This equation can be written in matrix form as $L(x \rightarrow y) = \mathbf{B}(x \rightarrow y)\mathbf{L}$, where \mathbf{L} is an infinite column matrix whose i -th entry is ℓ_i . When no

confusion can arise, we suppress the arguments and simply write $L = \mathbf{BL}$.

In the original formulation of radiosity, piecewise-constant functions were used as a basis [12]. In subsequent work on radiosity, Zatz [26] and Troutman and Max [24] used orthogonal polynomials, and Gortler *et al.* [13] used wavelets. In the more general context of radiance, the distribution of light leaving a patch has both spatial and angular variation. Immel *et al.* [15] used piecewise-constant basis functions for both spatial and angular variation. Later, Sillion *et al.* [21] used spherical harmonics for the angular variation and piecewise-constant basis functions for the spatial variation. In Sect. 3 we motivate and introduce our choice of basis, a wavelet basis for both spatial and angular variation.

Regardless of the choice of basis functions, we can obtain a system of equations for the unknown entries of \mathbf{L} by substituting $L = \mathbf{BL}$ and $L_e = \mathbf{BL}_e$ into the transport equation (2), and using linearity of the operator \mathcal{T} to yield

$$\mathbf{BL} = \mathbf{BL}_e + \mathcal{T}(\mathbf{BL}) = \mathbf{BL}_e + (\mathcal{T}\mathbf{B})\mathbf{L}.$$

Let $\langle f | g \rangle$ denote the standard inner product, $\int_{xy} f(x \rightarrow y) g(x \rightarrow y) dx dy$. If $\mathbf{F} = (f_1, f_2, \dots)$ and $\mathbf{G} = (g_1, g_2, \dots)$ are two row matrices of functions, let $[(\mathbf{F} | \mathbf{G})]$ be the matrix whose ij -th entry is $\langle f_i | g_j \rangle$. For clarity, we assume an orthonormal basis throughout this paper (introduction of dual basis functions [7, 19] is necessary for non-orthonormal bases). By applying the linear operator $[(\mathbf{B} | \cdot)]$ to both sides of the equation above and using orthonormality of the basis functions, we arrive at the infinite system of linear equations

$$\mathbf{L} = \mathbf{L}_e + \mathbf{TL}, \quad (3)$$

where $\mathbf{T} = [(\mathbf{B} | \mathcal{T}\mathbf{B})]$ is an infinite matrix representing the transport operator \mathcal{T} . The rs -th entry of \mathbf{T} is a *transport coefficient*, representing the influence of the coefficient of b_s on the coefficient of b_r . It can be written explicitly as

$$T_{r \leftarrow s} = \langle \bar{b}_r | \mathcal{T}b_s \rangle = \int_{xyz} \bar{b}_r(y \rightarrow z) f_r(x, y, z) G(x, y) b_s(x \rightarrow y) dz dy dx, \quad (4)$$

where the notation $r \leftarrow s$ is to emphasize that $T_{r \leftarrow s}$ represents the influence of the *sender* s on the *receiver* r .

3 A Wavelet Basis for Radiance

In this section we construct a basis for efficiently representing radiance distributions. Recent results by Beylkin *et al.* [4, 5], Alpert [1], Gortler *et al.* [13], Hanrahan *et al.* [14] and others indicate that significant performance gains can be achieved using a multiresolution basis. We first present some background on multiresolution analysis, and then describe one-dimensional wavelet bases and how they can be extended to four-dimensional bases for radiance distributions.

3.1 Multiresolution Analysis

Multiresolution analysis as formulated by Mallat [16] provides a convenient framework for studying multiresolution bases. There are two basic ingredients for a multiresolution analysis: an infinite chain of nested linear function spaces $V^0 \subset V^1 \subset V^2 \subset \dots$ and an inner product $\langle f | g \rangle$ defined on any pair of functions $f, g \in V^j$. The space V^j contains functions of resolution j , with resolution increasing as j increases. *Scaling functions* refer to bases for the spaces V^j . A function can be approximated by a sum of scaling functions.

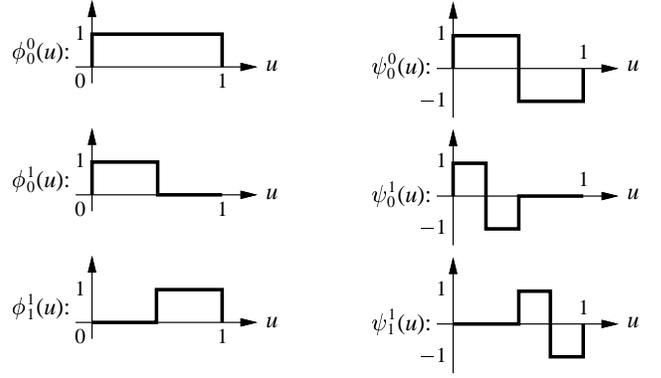


Figure 1 Some box functions $\phi_i^j(u)$ and Haar wavelets $\psi_i^j(u)$.

Alternatively, we can represent the same approximation as coarse scaling functions in V^0 along with detail at finer and finer resolutions. Detail is represented by functions in the *orthogonal complement spaces* W^j defined by

$$W^j \equiv \{f \in V^{j+1} \mid \langle f | g \rangle = 0 \ \forall g \in V^j\}.$$

Wavelets refer to bases for the orthogonal complement spaces W^j ; the spaces W^j are therefore called *wavelet spaces*.

Orthogonal complements are often written as $V^{j+1} = V^j \oplus W^j$ since, intuitively, wavelet space W^j includes the functions that are in V^{j+1} but “missing” from V^j . More formally, any function $f^{j+1} \in V^{j+1}$ can be written uniquely as an orthogonal decomposition $f^{j+1} = f^j + f_\perp^j$, where $f^j \in V^j$ and $f_\perp^j \in W^j$. The space V^j can be fully decomposed as

$$V^j = V^0 \oplus W^0 \oplus \dots \oplus W^{j-1}.$$

A *multiresolution basis* for V^j can be formed by selecting a scaling function basis for V^0 and wavelet bases for the spaces W^0, \dots, W^{j-1} . The scaling functions spanning V^0 represent coarse variation, while the wavelets provide detail at increasing resolutions.

For a more complete introduction to wavelets and their applications in computer graphics, see Stollnitz *et al.* [23].

3.2 Choice of Wavelet Basis

The simplest multiresolution basis in one dimension is the Haar basis [13]. The space V^j consists of piecewise-constant functions on $[0, 1]$ with discontinuities at $\{0, 1/2^j, 2/2^j, \dots, 1\}$. The space V^j is spanned by the Haar scaling functions $\phi_i^j(u)$, while the wavelet space W^j is spanned by piecewise-constant wavelets $\psi_i^j(u)$. A few Haar scaling functions and wavelets are shown in Fig. 1.

There are many alternatives to the Haar basis, each with advantages and disadvantages. One requirement for an efficient adaptive algorithm is the availability of fast quadrature formulas for the scaling functions and wavelets (and their duals, if non-orthonormal bases are used). In addition, due to the high dimensionality of the radiance transport problem, it is advantageous to have only one scaling function in space V^0 : a single scaling function leads to a single interaction between two patches at the coarsest level, while (as shown in Sect. 3.3) having two one-dimensional scaling functions leads to 16 four-dimensional scaling functions, requiring 256 interactions between two patches at the coarsest level. Finally, bounded-interval wavelets are preferable to wavelets with unbounded support, since it is unclear how radiance distributions would be artificially extended beyond the geometric extent of surface patches.

Among the wavelet bases that have the advantages outlined above, there are both continuous and discontinuous choices. There are cur-



Figure 2 Mapping the hemisphere to the unit square: we use gnomonic projection (a), followed by a radial “stretch” (b).

rently two families of bounded-interval continuous wavelets available in the literature: Daubechies wavelets adapted to the bounded interval [9], and bounded-interval B-spline wavelets [8]. Note that having continuous basis functions on each patch is not sufficient to ensure a continuous solution: continuity must also be enforced across boundaries of adjacent patches, or else the basis functions must be defined over complex shapes with arbitrary topology (for example, the complex floor shape in Fig. 4 in the color section).

We have experimented with bounded-interval B-spline wavelets [7], Daubechies wavelets, and the Haar basis. Of these, the Haar basis has many advantages, including orthogonality, compact support, and simple quadrature formulas. Although flatlets [13] have more vanishing moments¹ than the Haar basis, flatlets have wider support, requiring costly quadrature formulas. Multiwavelets [13] are constructed from higher-order polynomials, which also require costly quadrature formulas. The main disadvantage of the Haar basis, its discontinuities, can be ameliorated by performing a final gathering step during rendering [18].

3.3 A Four-Dimensional Wavelet Basis

Four-dimensional basis functions are required for representing radiance distributions: two variables describe spatial variation across a surface, and two variables describe angular variation. As is common, we split the surfaces into patches such that the spatial variables on each patch can be parameterized on the unit square $[0, 1]^2$. The domain of the radiance distributions is then $[0, 1]^2 \times H^2$, where H^2 is the unit hemisphere. By mapping H^2 onto $[0, 1]^2$, we can use tensor products of one-dimensional basis functions for angular variations, just as we do for spatial variations.

We use *gnomonic projection* to map between points in H^2 and points on a disc with radius $\pi/2$. As shown in Fig. 2(a), gnomonic projection maps great circles through the pole of H^2 to radial lines, and preserves arc length along these curves. We use this map because it is easily computed and introduces only mild distortion. This projection is followed by a radial “stretch” of the disc to exactly cover the unit square, as shown in Fig. 2(b). The composition of these mappings is a continuous and invertible mapping between H^2 and the unit square.

Sillion *et al.* [21] use spherical harmonics as basis functions for angular variations in radiance. These functions have the advantage of being naturally defined on the sphere, thereby eliminating the need for projection to the plane. The number of basis functions to represent a directional radiance distribution with spherical harmonics is comparable to the number needed in the Haar wavelet basis, so there is no immediate advantage of using the Haar wavelet basis. However, the Haar wavelets have local support while spherical harmonics have global support. Therefore, the transport matrix is dense for a spherical harmonics basis but sparse for a Haar wavelet basis.

We use the “nonstandard” wavelet basis [4], constructed from tensor products of univariate basis functions as follows. Let $\mathbf{u} = (u_1, u_2, u_3, u_4)$ denote a point in $[0, 1]^4$, and let $\mathbf{i} = (i_1, i_2, i_3, i_4)$ denote a 4-component multi-index of integers. The four-dimensional

scaling functions for V^j take the form

$$\phi\phi\phi\phi_{\mathbf{i}}^j(\mathbf{u}) \equiv \phi_{i_1}^j(u_1)\phi_{i_2}^j(u_2)\phi_{i_3}^j(u_3)\phi_{i_4}^j(u_4).$$

That is, the scaling functions for resolution j consist of all possible products of the one-dimensional scaling functions for resolution j . The four-dimensional wavelets spanning the orthogonal complement W^j are formed by taking all other products of scaling functions and wavelets for resolution j . These wavelets consist of 15 types:

$$\phi\phi\phi\psi_{\mathbf{i}}^j(\mathbf{u}), \phi\phi\psi\phi_{\mathbf{i}}^j(\mathbf{u}), \phi\psi\psi\phi_{\mathbf{i}}^j(\mathbf{u}), \dots, \psi\psi\psi\psi_{\mathbf{i}}^j(\mathbf{u}).$$

We take as our basis \mathbf{B} the set of basis functions spanning V^0, W^0, W^1, \dots for each patch in the scene.

4 A Wavelet Radiance Algorithm

We now turn to our wavelet-based radiance solution method. In some respects, our algorithm is similar to the approach described by Gortler *et al.* [13] for wavelet radiosity. However, there are a number of ways—in addition to the higher dimensionality—in which our algorithm differs significantly from previous radiosity work. Most significantly, our use of the standard operator decomposition eliminates the need for pushing and pulling procedures, and permits refinement of links at either end. We also describe how our refinement oracle serves to drive an adaptive quadrature scheme.

4.1 Main Algorithm

Initially, we project L_e into space V^0 , the space spanned by the scaling functions, to give $\tilde{\mathbf{L}}_e$. We also compute (as described in Sect. 4.2) the entries of $\tilde{\mathbf{T}}$ corresponding to interactions of scaling functions in V^0 with one another, giving $\tilde{\mathbf{T}}$. Quantities with a tilde are approximate, both because they represent truncated versions of infinite matrices and because they are computed numerically.

The main part of the algorithm alternates between computing an approximate radiance solution $\tilde{\mathbf{L}}$ and improving the finite representation of the transport operator $\tilde{\mathbf{T}}$:

```

procedure WaveletRadiance( $\tilde{\mathbf{T}}, \tilde{\mathbf{L}}_e$ ):
   $\tilde{\mathbf{L}} \leftarrow \tilde{\mathbf{L}}_e$ 
  repeat
     $\tilde{\mathbf{L}} \leftarrow \text{GaussSeidel}(\tilde{\mathbf{T}}, \tilde{\mathbf{L}}, \tilde{\mathbf{L}}_e)$ 
     $\tilde{\mathbf{T}} \leftarrow \text{Refine}(\tilde{\mathbf{T}}, \tilde{\mathbf{L}})$ 
  until visual convergence of  $\tilde{\mathbf{L}}$ 
end procedure

```

The main loop iterates until *visual convergence* is achieved, that is, until further refinement does not change the computed image significantly. We use Gauss-Seidel iteration to solve an approximate version of the discrete transport equation (3) given by

$$(\mathbf{I} - \tilde{\mathbf{T}})\tilde{\mathbf{L}} = \tilde{\mathbf{L}}_e.$$

The main algorithm calls on a *refinement oracle*, described in Sect. 4.3, to refine the radiance transport matrix.

Just as in previous hierarchical radiosity algorithms [10], the matrices $\tilde{\mathbf{T}}, \tilde{\mathbf{L}}$, and $\tilde{\mathbf{L}}_e$ are never formed explicitly. Entries of $\tilde{\mathbf{L}}$ and $\tilde{\mathbf{L}}_e$ are associated with the surface patches, while entries of $\tilde{\mathbf{T}}$ are stored as “links” between radiance coefficients.

Note that the algorithm presented by Gortler *et al.* [13] requires “Push” and “Pull” procedures to distribute transported radiosity among the levels of a hierarchy between Gauss-Seidel iterations. By using the standard operator decomposition and representing radiance in terms of wavelets rather than scaling functions, our algorithm eliminates the pushing and pulling procedures. On the

¹As discussed by Alpert [1] and Gortler *et al.* [13], for a smooth operator, an increased number of vanishing moments will increase the sparsity of the discrete approximation to that operator.

other hand, the nonstandard operator decomposition is in theory more sparse than the standard decomposition for an operator that is smooth. It is not clear whether or not this theoretical advantage of the nonstandard decomposition has a practical implication for a piecewise-smooth operator like the light transport operator.

4.2 Computing Transport Coefficients

Each transport coefficient $T_{r \leftarrow s}$ is defined in Equation (4) as an inner product that results in a six-dimensional integral. For example, the influence of wavelet $\psi\phi\psi\phi_{\mathbf{i}_s}(\mathbf{u}_s)$ on wavelet $\psi\phi\psi\phi_{\mathbf{i}_r}(\mathbf{u}_r)$ is $T_{r \leftarrow s} = \langle \psi\phi\psi\phi_{\mathbf{i}_r} | \mathcal{T}\psi\phi\psi\phi_{\mathbf{i}_s} \rangle$. If we write $\mathbf{u} = (x, \omega)$, where $x = (u_1, u_2)$ denotes spatial components and $\omega = (u_3, u_4)$ denotes angular components, then the inner product takes the form

$$\begin{aligned} T_{r \leftarrow s} &= \langle \psi\phi\psi\phi_{\mathbf{i}_r} | \mathcal{T}\psi\phi\psi\phi_{\mathbf{i}_s} \rangle & (5) \\ &= \int \psi\phi\psi\phi_{\mathbf{i}_r}(x_r, \omega_r) f_r(x_s, x_r, \omega_r) G(x_s, x_r) \psi\phi\psi\phi_{\mathbf{i}_s}(x_s, \omega_s) d\omega_r dx_r dx_s \\ &= \int \left[\int \psi\phi\psi\phi_{\mathbf{i}_r}(x_r, \omega_r) f_r(x_s, x_r, \omega_r) d\omega_r \right] \\ &\quad \cdot G(x_s, x_r) \psi\phi\psi\phi_{\mathbf{i}_s}(x_s, \omega_s) dx_r dx_s. \end{aligned}$$

Here ω_s is considered to be a function of x_s and x_r , since the direction at the sender must lie along the line between sending and receiving positions. Note that only the BRDF and the receiving basis function depend on ω_r . Our numerical integration routine samples these two functions in its innermost loop, while the remaining functions are evaluated only as the positional variables change.

We approximate integrals such as the one above using (slightly jittered) uniform sampling of the integrand. More accurate rules such as Gauss-Legendre or Gauss-Kronrod quadrature could also be used [7, 13, 26].

4.3 Refinement

In many applications of wavelets in numerical analysis, the goal is to obtain a sparse representation of a given matrix, thereby making repeated matrix–vector multiplications much faster [4]. In wavelet-based approaches to global illumination, the cost of explicitly constructing an entire transport matrix far outweighs the expense of any matrix–vector multiplications that follow. Therefore, it is essential to restrict the number of computed transport coefficients.

The goal of the refinement oracle is to determine where to refine $\tilde{\mathbf{T}}$ to better approximate \mathbf{T} . The two most important sources of error are:

- *truncation error* due to significant entries missing from $\tilde{\mathbf{T}}$, and
- *quadrature error* in computing the entries of $\tilde{\mathbf{T}}$.

In this section we describe how our oracle reduces truncation error. Section 4.4 outlines a method for simultaneously reducing quadrature errors.

Our refinement oracle is a generalization of the brightness refinement criterion for hierarchical radiosity [14] and the oracle used by Gortler *et al.* for wavelet radiosity [13]. The idea is to estimate the amount of light that would be transported if a new transport coefficient were to be added to $\tilde{\mathbf{T}}$. If this quantity falls below some threshold, then it is likely that the expensive computation of the transport coefficient can be avoided without resulting in significant error in the solution.

For a given link λ between a pair of sending and receiving basis functions, we consider refining both at the sending end and at the

receiving end of λ . In either case, we multiply a sending basis function coefficient by an estimate of the transport coefficient for the link λ_{new} under consideration. In our implementation, the transport coefficient for λ_{new} is estimated by the variation in the kernel evaluations for λ . (This variation is stored along with the transport coefficient on link λ .) By contrast, Gortler *et al.* use a polynomial interpolant rather than sample variation to estimate kernel smoothness.

Suppose the oracle decides to refine the receiving end of a link. Then new links are created as follows: When refining a link to a scaling function in space V^0 , links to all 15 wavelets in space W^0 are created. When refining a link to a wavelet in space W^j , links to the overlapping wavelets of the same type in W^{j+1} are created. (In the case of the Haar basis, $2^4 = 16$ new links are created.) A similar process occurs when the oracle decides refinement is needed at the sending end of a link.

Note that our refinement procedure can refine each end of a link independently, and in our algorithm links are never destroyed. By contrast, the approach described by Gortler *et al.* removes a link at one level of the hierarchy and replaces it with multiple links at a finer level of detail, thereby refining both ends simultaneously.

4.4 Adaptive Quadratures

If we always use a numerical integration rule of high accuracy to compute transport coefficients, computation may be wasted evaluating the kernel for many interactions that have little effect on the final image. On the other hand, the significant coefficients have to be computed accurately; otherwise, the solution will not converge to the correct value. It is therefore advantageous to use an adaptive integration technique that reduces error only for significant transport coefficients.

For time efficiency, we would like to store the values of all kernel evaluations that have already been computed in order to reuse them for improved quadratures. Unfortunately, space limitations prohibit this approach. Instead, we store only the evaluations from the most recently computed transport coefficient. These new kernel evaluations can be used to update any transport coefficient whose sending and receiving basis functions overlap with those of the new transport coefficient.

When a transport coefficient $T_{r \leftarrow s}$ is refined and a new entry $T_{r' \leftarrow s'}$ is computed, we check whether or not a transport coefficient between basis functions with the same supports as r' and s' has already been computed. If so, we need do no more; if not, the samples used to compute $T_{r' \leftarrow s'}$ are used to update $T_{r \leftarrow s}$. Updating a transport coefficient is done in two steps. First, the original samples of the integrand within the supports of r' and s' are recomputed and subtracted from $T_{r \leftarrow s}$. Next, the new kernel evaluations are multiplied by the appropriate basis functions, weighted by area, and added to $T_{r \leftarrow s}$. This approach to adaptive quadrature is less expensive than simply recomputing $T_{r \leftarrow s}$ more accurately upon refinement, since we reuse the costly kernel samples whenever possible.

Although we have only implemented this adaptive integration technique for the Haar basis, it could be extended to other wavelet bases. However, it is not immediately apparent how this approach would generalize to integration rules using nonuniform sampling.

5 Implementation features

In this section, we describe features of our implementation. Objects in the scene can consist of flat quadrilaterals and tensor-product Bézier patches, and can have isotropic or anisotropic reflection. The light sources can have spatial and angular variation. We use importance to restrict refinement to the light transports that influence the

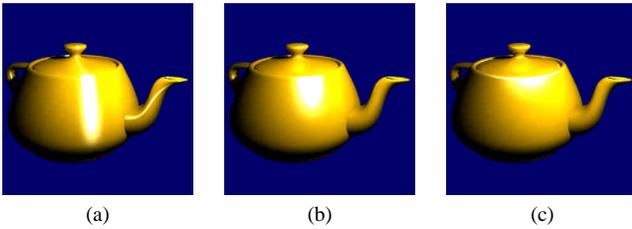


Figure 3 Ward’s reflection model: (a) anisotropic reflection with $\alpha_u = 0.1, \alpha_v = 0.5$; (b) isotropic reflection with $\alpha_u = \alpha_v = 0.2$; (c) anisotropic reflection with $\alpha_u = 0.5, \alpha_v = 0.1$.

final image the most, and employ a final gathering step to improve the visual quality of the solution.

5.1 Surface Geometry

Any parametric surface representation can be used by our algorithm, so long as we are able to compute a position, surface normal, and differential area associated with a given parametric point (u_1, u_2) , and determine the intersection of a ray with the surface. Our implementation currently handles tensor-product Bézier patches and quadrilaterals. It would also be straightforward to add nonuniform rational B-spline surfaces. The images in Fig. 3 show a teapot consisting of 28 Bézier patches.

5.2 Reflection Models and Texture Maps

We use the Ward isotropic and anisotropic reflection models [25] since they are physically valid and fast to evaluate. Examples of this reflection model can be seen in Fig. 3. In addition, we use spatially varying reflectances to simulate details of the materials in the scene. We take the BRDF to be the product of a spatially-varying texture and the angular variation of the Ward model. Figure 4 demonstrates both texture-mapped and anisotropic reflectance functions.

In the course of numerically approximating a transport coefficient, the geometric term and the BRDF are sampled at a number of points. The reflectance for each point is determined by a look-up in a texture map, multiplied by the angular variation given by Ward’s model. Gershbein *et al.* [11] present an alternative approach, using wavelet decompositions of textures for radiosity.

5.3 Light Sources

By storing the wavelet decomposition of an image as the initial coefficients on a patch, we can model a light source that emits a spatially-varying radiance (like a television screen). In general, not all coefficients of the emitting image will have links from them, but the coefficients are ready to be transported into the scene if the refinement procedure so decides. This technique allows a complex environment to be displayed using simple geometry.

A simple approach to angular variation is to let the emission depend upon direction. For example, we model “spotlights” using a Phong-like function, in which emission is proportional to some power of the cosine of the angle between the emission direction and the surface normal of the patch. The spotlights appear dark from most directions because of the very narrow distribution of light they emit.

We demonstrate the use of spotlights and a spatially-varying emitter (the outdoor environment seen through the window) in Fig. 4. More complex effects such as a slide projector or sunlight through a stained-glass window could be modeled by combining spatial and angular variations in an emitter.

5.4 Importance

In order to maintain a tractably small problem for complex scenes, we use importance-driven refinement for a view-dependent solution. Importance was described in Smits *et al.* [22] for radiosity and in Christensen *et al.* [6], Aupperle and Hanrahan [2], and Patanaik [17] for radiance. Briefly, importance measures the fraction of light leaving a point that will reach the eye.

The algorithm combines estimates of importance and radiance to drive the global solution, allowing it to exploit view-dependent information as part of an adaptive refinement scheme. We use exitant directional importance [6], since it satisfies the same transport equation as radiance. Importance can therefore be discretized in the same manner as radiance and transported by the same links. The only difference between importance and radiance is that radiance is emitted by light sources, while exitant directional importance is emitted by the eye.

Smits *et al.* [22] showed that importance gives a substantial speed-up for a complex diffuse scene. For glossy reflections, the gain in speed is even greater, due to the directionality of radiance and importance: a directional interaction is refined only if the amount of transported radiance *in that direction* is both large and important. Note that we can get arbitrarily large speed-ups, compared to a solution obtained without using importance, by choosing a sufficiently complex scene where many parts do not contribute significantly to the final image.

The first three images of Fig. 4 show a complex scene viewed from above. The radiance emitted by the spotlights and reflected in the scene is shown in Fig. 4(a). Importance is emitted from the eye and reflected to the important parts of the scene, as shown in Fig. 4(b). This picture demonstrates how small a fraction of the model significantly influences the solution visible from the eye. Figure 4(c) is a gray-scale encoding of the number of links between the basis functions on each surface patch. This “refinement image” verifies that most work is performed in areas that are both bright and important.

5.5 Final Gather

Following the ideas that Reichert [18] used for radiosity, we have implemented a final radiance gathering step. For each pixel in the image, we perform a final gathering of light to the surface point that corresponds to the midpoint of the pixel. For each sending basis function, we evaluate a simplified version of the integral in Equation (5). Since the receiving position is fixed and the radiance is reflected towards the eye, the integration is over only sending positions.

Formally, this final gather corresponds to changing to a piecewise-constant basis, where the support of each basis function is the projection of a pixel onto a surface in the scene. This basis is tailored to be visually pleasing. The final gather smooths the discontinuities in the wavelet representation and makes highlights, textures, and shadows crisper. The improvement brought about by the final gather can be seen by comparing Figs. 4(d) and 4(e).

Another way of thinking about the final gathering step is in the context of distribution ray tracing. When a ray emanating from the eye intersects a surface in the scene, a group of reflected rays are traced from the intersection point to points on other surfaces in the scene. A constant number of rays are cast to the support of each basis function in the radiance solution, so the directions of the rays are guided by the radiance solution. Thus, the most refined areas of the radiance solution are sampled the most by the distribution of reflected rays. Note that the costly “explosion” of the number of recursive bounces used in distribution ray tracing is avoided, and that the final gather requires no additional memory.

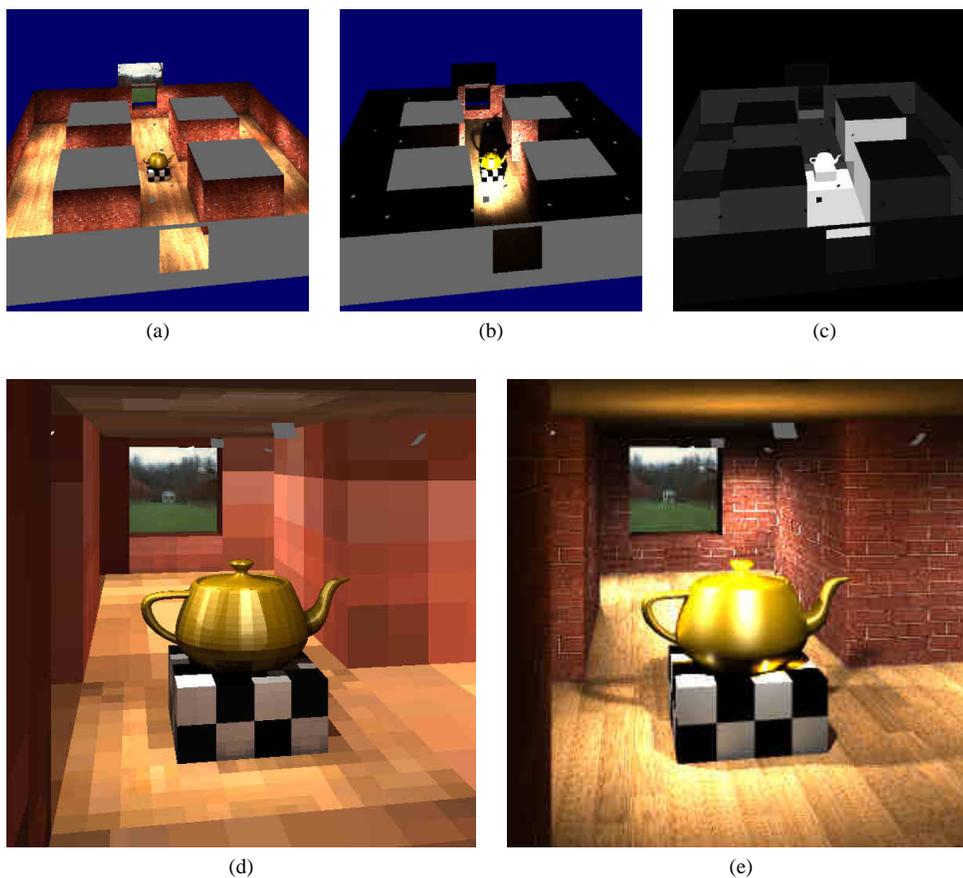


Figure 4 Solutions for a complex scene: (a) radiance seen from above; (b) importance seen from above; (c) gray-scale representation of refinement; (d) radiance solution without final gather; (e) radiance solution with final gather.

6 Results

As a test scene, we used a maze of hallways with a glossy Bézier-patch teapot in the center (see Fig. 4). The scene consists of 152 patches, including 28 Bézier patches, and has 8,802 mutually visible pairs of patches. The teapot’s reflectance function is anisotropic with specularities $\alpha_u = 0.2$ and $\alpha_v = 0.5$, specular reflectivity $\rho_s = (0.1, 0.1, 0.1)$, and diffuse reflectivity $\rho_d = (0.2, 0.15, 0)$. The illumination consists of 24 “spotlights,” patches that emit directional radiance. There is a patch outside the window that emits light according to a scanned image of an outdoor scene, giving the appearance of a full environment beyond the window. There is also a small patch representing the eye in the hallway in front of the teapot. The eye patch emits importance in the direction of the teapot, just like a spotlight emits light.

The program uses the four-dimensional wavelet basis described in Sect. 3.3 and begins by creating 8,802 links between scaling functions in V^0 . In the first refinement, 7,491 new links are created, and in the next refinement 123,560 new links are created. Running times on a DEC Alpha machine were approximately five minutes to compute the initial transport coefficients between scaling functions in V^0 , then 110 minutes to iterate the main algorithm and refine as far as V^4 in important parts of the scene, and 15 minutes to render a 600×600 image using ray casting and evaluation of the solution. The final gather takes another two hours, so the time for the final gather is comparable to the computation-time for the solution.

Note the interreflections: there is significant color bleeding from the brightly illuminated teapot to the dim ceiling, and the white squares on the pedestal are brightly reflected in the bottom of the teapot.

7 Conclusion

We have presented an efficient method for simulating light transport in an environment with diffuse and glossy reflections. We use wavelet basis functions to represent the four-dimensional radiance distribution associated with surfaces in a scene. Wavelets adapt to the solution, so in areas with little spatial or angular variation a coarse solution is computed, and in areas with greater detail a more refined solution is found.

In contrast to previous algorithms for wavelet radiosity, we use a standard decomposition of the transport operator, and we represent radiance as a weighted sum of wavelets rather than scaling functions. We do not use pushing and pulling procedures, and we are able to refine just that end of a link for which the estimated improvement is greatest. In order to obtain accurate numerical integration without the expense of extraneous samples, we have also developed adaptive integration rules for the transport coefficients.

Radiance transport is formulated as a multidimensional Fredholm integral equation of the second kind. Thus, our approach may benefit other fields in which such equations arise—numerical analysis, finite element analysis, and particle transport simulation, for example.

There are a number of areas in which we foresee future work. A comparison of wavelet bases for radiance should examine rates of convergence, quadrature expense and accuracy, continuity properties, and the amount of work required to obtain a solution of a given accuracy. Standard and nonstandard operator decomposition should be compared for a piecewise-smooth kernel typical of global il-

lumination with partial occlusion. It would also be interesting to compare wavelets for two-point transport with a three-point transport algorithm, as described by Aupperle and Hanrahan [3, 2] for piecewise-constant basis functions. Finally, an investigation of the final gathering step should determine whether or not it improves the numerical accuracy of the solution in addition to improving its visual appearance.

There are many possible extensions to the present algorithm. Wavelet bases are not suited to the representation of ideal specular reflections. Instead, a ray tracing step for ideal specular reflection could be incorporated in the same fashion as in Sillion *et al.* [21]. Furthermore, surfaces that transmit light in addition to reflecting it could be incorporated into our algorithm by using a wavelet basis defined for the entire sphere of directions.

Acknowledgements

This work was supported by NSF Presidential and National Young Investigator awards (CCR-8957323 and CCR-9357790), an NSF Graduate Research Fellowship, the University of Washington Graduate Research Fund (award 75-1721), the Technical University of Denmark, and the Danish Research Council.

References

- [1] Bradley K. Alpert. Sparse representations of smooth linear operators. Ph.D. thesis, Yale University, 1990.
- [2] Larry Aupperle and Pat Hanrahan. Importance and discrete three point transport. In *Proceedings of the Fourth Eurographics Workshop on Rendering* (Paris, June 1993), pages 85–94.
- [3] Larry Aupperle and Pat Hanrahan. A hierarchical illumination algorithm for surfaces with glossy reflection. In *Proceedings of SIGGRAPH '93*, pages 155–162. ACM, New York, 1993.
- [4] G. Beylkin, R. Coifman, and V. Rokhlin. Fast wavelet transforms and numerical algorithms I. *Communications on Pure and Applied Mathematics*, 44(2):141–183, March 1991.
- [5] G. Beylkin, R. R. Coifman, and V. Rokhlin. Wavelets in numerical analysis. In Mary Beth Rushkai *et al.*, editors, *Wavelets and Their Applications*, pages 181–210. Jones and Bartlett, Boston, 1992.
- [6] Per H. Christensen, David H. Salesin, and Tony D. DeRose. A continuous adjoint formulation for radiance transport. In *Proceedings of the Fourth Eurographics Workshop on Rendering* (Paris, June 1993), pages 95–104.
- [7] Per H. Christensen, Eric J. Stollnitz, David H. Salesin, and Tony D. DeRose. Importance-driven wavelet radiance. Technical Report 94-01-05, Department of Computer Science and Engineering, University of Washington, January 1994.
- [8] Charles K. Chui and Ewald Quak. Wavelets on a bounded interval. In D. Braess and L. L. Schumaker, editors, *Numerical Methods in Approximation Theory*, volume 9, pages 53–75. Birkhauser Verlag, Basel, 1992.
- [9] A. Cohen, I. Daubechies, and J. C. Feauveau. Biorthogonal bases of compactly supported wavelets. *Communications on Pure and Applied Mathematics*, 45(5):485–500, June 1992.
- [10] Michael F. Cohen and John R. Wallace. *Radiosity and Realistic Image Synthesis*. Academic Press Professional, Cambridge, MA, 1993.
- [11] Reid Gershbein, Peter Schröder, and Pat Hanrahan. Textures and radiosity: Controlling emission and reflection with texture maps. In *Proceedings of SIGGRAPH '94*, pages 51–58. ACM, New York, 1994.
- [12] Cindy M. Goral, Kenneth E. Torrance, Donald P. Greenberg, and Bennett Battaile. Modeling the interaction of light between diffuse surfaces. In *Proceedings of SIGGRAPH '84*, pages 213–222. ACM, New York, 1984.
- [13] Steven J. Gortler, Peter Schröder, Michael F. Cohen, and Pat Hanrahan. Wavelet radiosity. In *Proceedings of SIGGRAPH '93*, pages 221–230. ACM, New York, 1993.
- [14] Pat Hanrahan, David Salzman, and Larry Aupperle. A rapid hierarchical radiosity algorithm. In *Proceedings of SIGGRAPH '91*, pages 197–206. ACM, New York, 1993.
- [15] David S. Immel, Michael F. Cohen, and Donald P. Greenberg. A radiosity method for non-diffuse environments. In *Proceedings of SIGGRAPH '86*, pages 133–142. ACM, New York, 1986.
- [16] Stephane Mallat. A theory for multiresolution signal decomposition: The wavelet representation. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 11(7):674–693, July 1989.
- [17] Sumanta N. Pattanaik. Computational methods for global illumination and visualisation of complex 3D environments. Ph.D. thesis, Birla Institute of Technology and Science, 1993.
- [18] Mark C. Reichert. A two-pass radiosity method driven by lights and viewer position. Master's thesis, Cornell University, 1992.
- [19] Peter Schröder, Steven J. Gortler, Michael F. Cohen, and Pat Hanrahan. Wavelet projections for radiosity. *Computer Graphics Forum*, 13(2):141–151, June 1994.
- [20] Min-Zhi Shao, Qun-Sheng Peng, and You-Dong Liang. A new radiosity approach by procedural refinements for realistic image synthesis. In *Proceedings of SIGGRAPH '88*, pages 93–102. ACM, New York, 1988.
- [21] François X. Sillion, James R. Arvo, Stephen H. Westin, and Donald P. Greenberg. A global illumination solution for general reflectance distributions. In *Proceedings of SIGGRAPH '91*, pages 187–196. ACM, New York, 1991.
- [22] Brian E. Smits, James R. Arvo, and David H. Salesin. An importance-driven radiosity algorithm. In *Proceedings of SIGGRAPH '92*, pages 273–282. ACM, New York, 1992.
- [23] Eric J. Stollnitz, Tony D. DeRose, and David H. Salesin. Wavelets for computer graphics: A primer. *IEEE Computer Graphics and Applications*, 15(3):76–84, May 1995 (part 1) and 15(4):75–85, July 1995 (part 2).
- [24] Roy Troutman and Nelson L. Max. Radiosity algorithms using higher order finite elements. In *Proceedings of SIGGRAPH '93*, pages 209–212. ACM, New York, 1993.
- [25] Gregory J. Ward. Measuring and modeling anisotropic reflection. In *Proceedings of SIGGRAPH '92*, pages 265–273. ACM, New York, 1992.
- [26] Harold R. Zatz. Galerkin radiosity. In *Proceedings of SIGGRAPH '93*, pages 213–220. ACM, New York, 1993.