A Framework for the Analysis of Error in Global Illumination Algorithms

James Arvo
Kenneth Torrance
Brian Smits

Program of Computer Graphics*
Cornell University

Abstract

In this paper we identify sources of error in global illumination algorithms and derive bounds for each distinct category. Errors arise from three sources: inaccuracies in the boundary data, discretization, and computation. Boundary data consist of surface geometry, reflectance functions, and emission functions, all of which may be perturbed by errors in measurement or simulation, or by simplifications made for computational efficiency. Discretization error is introduced by replacing the continuous radiative transfer equation with a finite-dimensional linear system, usually by means of boundary elements and a corresponding projection method. Finally, computational errors perturb the finite-dimensional linear system through imprecise form factors, inner products, visibility, etc., as well as by halting iterative solvers after a finite number of steps. Using the error taxonomy introduced in the paper we examine existing global illumination algorithms and suggest new avenues of research.

CR Categories and Subject Descriptors: I.3.7 [Computer Graphics]: Three-Dimensional Graphics and Realism.

Additional Key Words and Phrases: boundary elements, discretization, error bounds, global illumination, linear operators, projection methods, radiosity, reflectance functions.

1 Introduction

The role of global illumination algorithms is to simulate the interaction of light with large-scale geometry for the purpose of image synthesis. The greatest challenges in this endeavor have been those of accuracy and efficiency. While issues of efficiency have been addressed frequently in computer graphics, error analysis has received comparatively little attention. Indeed, the subject is difficult to even approach for several reasons. First, there is no universally accepted definition of accuracy for global illumination. Although accuracy connotes a quantitative comparison with some ideal, the ideal may be either empirical or theoretical, and the comparison may assume a range of forms from purely mathematical to perceptual. A second obstacle has been a lack of mathematical formalisms for global error analysis. Thus far, analysis has been confined to specific sub-problems, and no systematic treatment of the subject has been presented.

In this paper we attempt to establish a framework for error analysis within a well-defined class of global illumination problems. Specifically, we address the problem of approximating solutions to a form of the rendering equation [18] given imprecise data for geometry, reflectance functions, and emission functions. We further assume that the approximation is to be assessed quantitatively by its distance from the theoretical solution. By employing radiometric quantities and physically-motivated measures of error we temporarily sidestep the difficult issues of display and perception. Moreover, in the present work we exclude participating media, transparent surfaces, and probabilistic solution methods. The problem domain that we address can be summarized as follows:

- Radiometric quantities (radiance, reflectivities, emission)
- A linear model of radiative transfer
- Direct radiative exchange among opaque surfaces
- Deterministic boundary element formulations

Given these restrictions, we derive error bounds in terms of potentially known quantities, such as bounds on emission, reflectivity, and measurement error. To obtain these error bounds, we draw upon the general theory of integral equations [2] as well as the more abstract theory of operator equations [1, 21].

While the analysis is carried out in an abstract setting, we shall not lose sight of practical considerations. For instance, in simulating a given physical environment, perhaps under varied lighting conditions, how accurately must the reflectance functions be measured? Or, when simulating radiant transfer among diffuse surfaces, how important is it to use higher-order elements? Can we expect higher accuracy by using analytic area-to-area instead of point-to-area form factors? Finally, how accurate must visibility computations be for global illumination? While we do not provide definitive answers to these questions, the analysis presented here introduces a formalism and a starting point for determining quantitative answers.

2 Radiative Transfer

Global illumination involves the processes of light emission, reflection, redistribution, shadowing and, ultimately, absorption in an environment. These are physical processes governed by the equations of radiative transfer. In general, radiative transfer equations apply at both microscopic and macroscopic scales and are based primarily on the first and second laws of thermodynamics. These laws respectively affirm that thermal energy is conserved and flows

*580 Engineering and Theory Center Building, Ithaca, New York 14853

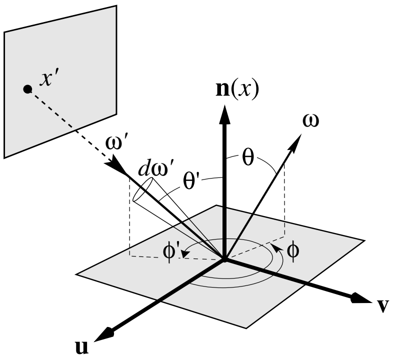

Figure 1: *The local coordinates of a bidirectional reflectance function. The incident vector ω' and the reflected vector ω are in world coordinates. The angles define a local coordinate system.*

irreversibly from regions of high potential (temperature) to regions of low potential [30]. The two laws also yield certain reciprocity conditions such as those exhibited by reflection functions and form factors.

In the present paper, we use a geometrical optics formalism to describe radiative transfer processes at the large scale; interference and diffraction are ignored at this scale. We restrict physical or wave optics effects to the level of scattering and emission at surfaces, where the radiative wavelength may be comparable in size to surface features. Further, monochromatic radiation is assumed, and we neglect energy transfer between wavelength bands (for instance, by absorption and re-emission at surfaces). Given the forgoing physical assumptions, we specify in this section the continuous governing equation for global illumination.

Let $\mathcal{M}$ denote the collection of all surfaces in an environment, which we assume to form an enclosure for simplicity. Let $\mathcal{X}$ be a space of real-valued functions defined on $\mathcal{M} \times \mathcal{S}^2$; that is, over all surface points and angular directions in the unit sphere $\mathcal{S}^2$. Given the *surface emission* function $g \in \mathcal{X}$, which specifies the origin and directional distribution of emitted light, we wish to determine the *surface radiance function* $f \in \mathcal{X}$ that satisfies

$$f(x,\omega) = g(x,\omega) + \int_{\Omega_i} k(x;\omega' \to \omega) f(x',\omega') \cos\theta' \, d\omega', \qquad (1)$$

where Ω_i is the hemisphere of incoming directions, k is a directional reflectivity, θ' is the angle of the incident beam contained in the solid angle $d\omega'$, and x' is a point on a distant surface determined by x and ω'. See Figure 1. Equation (1) is expressed in radiometric variables but otherwise embodies the same physical principles as the rendering equation introduced by Kajiya [18]. The radiometric formulation is essential for the present error analysis.

The function k appearing in equation (1) is the bidirectional reflectance distribution function at each surface point phrased in terms of direction vectors:

$$k(x;\omega' \to \omega) \equiv \rho_x(\theta',\phi',\theta,\phi) \qquad (2)$$

where (θ',ϕ') and (θ,ϕ) are the polar angles of the incident and reflected directions respectively. The function k has two crucial properties that follow from the thermodynamic principles of energy conservation and reciprocity. By conservation we have that

$$\int_{\Omega_o} k(x;\omega' \to \omega) \, \cos\theta \, d\omega \le 1, \qquad (3)$$

where $\omega' \in \Omega_i$ and Ω_o is the outgoing hemisphere. Equation (3) states that the energy reflected from a surface cannot exceed that of

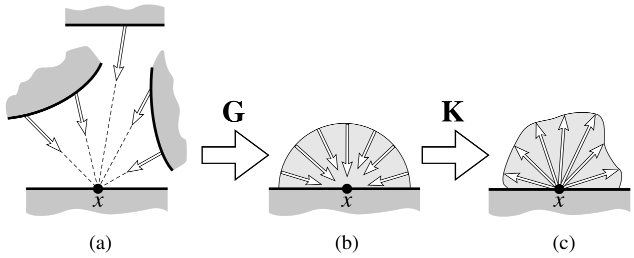

Figure 2: *The actions of **G** and **K** at a single point x. The operator **G** converts (a) distant surface radiance directed toward x into (b) local field radiance, where (c) it is again mapped into surface radiance by **K**.*

the incident beam. Reciprocity states that

$$k(x;\omega' \to \omega) = k(x;\omega \to \omega') \qquad (4)$$

for all $\omega' \in \Omega_i$ and $\omega \in \Omega_o$. These facts play a major role in the following analysis.

Note that the implicit function x' in equation (1) is the means of constructing the distribution of energy impinging on a surface from the distribution of energy leaving distant surfaces; that is, it constructs local *field radiance* from distant *surface radiance*. The connection afforded by $x'(x,\omega')$ corresponds to the intuitive operation of tracing a ray from x in the direction $-\omega'$. This simple coupling is a consequence of steady-state radiance being invariant along rays in free space; in the presence of participating media the coupling is replaced by the equation of transfer along the ray [6].

The balance equation for direct radiative exchange between surfaces was perhaps first presented in the form of equation (1) by Polyak [31]; in current radiative heat transfer literature this form is used when participating media are ignored [27]. Note that form factors do not appear directly in the integral form of the balance equation, but result from the $\cos\theta' \, d\omega'$ term for certain approximations.

2.1 Linear Operators

Integral equations such as equation (1) belong to the more general class of *operator equations*. Operator equations tend to be more concise than their integral counterparts while capturing algebraic properties essential for error analysis. Operator equations were first applied to global illumination by Kajiya [18] although their connection with integral equations of a similar nature has been profitably explored for nearly a century [4].

Equation (1) can be expressed as an operator equation in numerous ways. We shall construct the central operator from two simpler ones, each based on standard radiometric concepts. The representation given here has two novel features: first, it cleanly separates notions of geometry and reflection into distinct linear operators, and secondly, it employs an integral operator with a cosine-weighted measure. Both of these features simplify the subsequent error analysis.

We first define the *local reflection operator* **K**, an integral operator accounting for the scattering of incident radiant energy at surfaces. This operator is most easily defined by showing its action on an arbitrary field radiance function h:

$$(\mathbf{K}h)(x,\omega) \equiv \int_{\Omega_i} k(x;\omega' \to \omega) \, h(x,\omega') \, d\mu(\omega'). \qquad (5)$$

The **K** operator maps the field radiance function h to the surface radiance function after one reflection (Figures 2b-c). The operation is *local* in that the transformation occurs at each surface point in isolation. In equation (5) we have adopted the notation of Lebesgue integration [32] to introduce a measure μ that incorporates the cosine weighting. The essential relationship with the differential solid angle $d\omega'$ used in equation (1) is

$$d\mu(\omega') \equiv \cos\theta'\, d\omega'. \tag{6}$$

The new notation hides the ubiquitous cosines and, more importantly, emphasizes that the cosine is an artifact of surface integration. This observation allows the kernel k to inherit the reciprocity property of reflectance functions and simplifies some of the following analysis.

Next, we define the *field radiance operator* $\mathbf{G}$, a linear operator that expresses the incident field radiance at each point in terms of the surface radiance of the surrounding environment (Figures 2a-b). Showing the action of $\mathbf{G}$ on a surface radiance function h, we have

$$(\mathbf{G}h)(x,\omega) \equiv h(x'(x,\omega),\omega). \tag{7}$$

The $\mathbf{G}$ operator expresses non-local point-to-point visibility operations as a mapping on the space of radiance functions. Defining $\mathbf{G}$ in this way allows us to factor out the implicit function $x'(x,\omega)$ from the integral in equation (1). It is easily verified that $\mathbf{KG}$ is equivalent to the original integral and that both $\mathbf{K}$ and $\mathbf{G}$ are linear. Therefore, we can write equation (1) as

$$f = g + \mathbf{KG}f, \tag{8}$$

which is a linear operator equation of the *second kind*.

To highlight the linear relationship between surface radiance and surface emission, equation (8) can be written more compactly as

$$\mathbf{M}f = g, \tag{9}$$

where the linear operator $\mathbf{M}$ is defined by

$$\mathbf{M} \equiv \mathbf{I} - \mathbf{KG}, \tag{10}$$

and $\mathbf{I}$ is the identity operator. We also note that $\mathbf{K}$ and $\mathbf{G}$ are self-adjoint in a natural sense, so the adjoint of $\mathbf{M}$ can be written

$$\mathbf{M}^* = \mathbf{I} - \mathbf{GK}, \tag{11}$$

which is useful in importance-driven global illumination.

A global illumination problem is essentially a 3-tuple $(\mathcal{M}, \mathbf{M}, g)$ of surfaces, linear operator, and emission function. In the following sections we explore numerical methods for solving equation (9) and investigate sources of error.

3 Boundary Elements and Projection Methods

By far the most common methods for solving global illumination problems are those employing surface discretizations, which are essentially *boundary element* methods [26]. In more abstract terms, boundary element methods are themselves *projection methods* whose role is to recast infinite-dimensional problems in finite dimensions. In this section we pose the problem of numerical approximation for global illumination in terms of projections. This level of abstraction will allow us to clearly identify and categorize all sources of error while avoiding implementation details.

The idea behind boundary element methods is to construct an approximate solution from a known finite-dimensional subspace $\mathcal{X}_n \subset \mathcal{X}$, where the parameter n typically denotes the dimension of the subspace. For global illumination the space $\mathcal{X}_n$ may consist of n boundary elements over which the radiance function is constant. Alternatively, it may consist of fewer boundary elements, but with internal degrees of freedom, such as tensor product polynomials [41], spherical harmonics [35], or wavelets [12]. In any case, each element of the function space $\mathcal{X}_n$ is a linear combination of a finite number of basis functions, $u_1, \ldots, u_n$. That is,

$$\mathcal{X}_n = \text{span}\{u_1, \ldots, u_n\}. \tag{12}$$

Given a set of basis functions, we seek an approximation f_n from the space $\mathcal{X}_n$ that is "close" to f in some sense. By virtue of the finite-dimensional space, finding f_n is equivalent to determining n unknown coefficients $\alpha_1, \ldots, \alpha_n$ such that

$$f_n = \sum_{j=1}^{n} \alpha_j u_j. \tag{13}$$

There are many possible methods for selecting such an approximation from $\mathcal{X}_n$, each motivated by a specific notion of closeness and the computational requirements of finding the approximation.

A universal feature of discrete boundary element approaches is that they operate using a finite amount of "information" gathered from the problem instance. For projection methods this is done in the following way. We select $f_n \in \mathcal{X}_n$ by imposing a finite number of conditions on the *residual error*, which is defined by

$$r_n \equiv \mathbf{M}f_n - g. \tag{14}$$

Specifically, we attempt to find f_n such that r_n simultaneously satisfies n linear constraints. Since we wish to make the residual "small", we set

$$\phi_i(r_n) = 0, \tag{15}$$

for $i = 1, 2, \ldots n$, where the $\phi_i : \mathcal{X} \to \mathbb{R}$ are linear functionals. The functionals and basis functions together define a projection operator, as we show in section 6.2.

Any collection of n linearly independent functionals defines an approximation f_n by "pinning down" the residual error with sufficiently many constraints to uniquely determine the coefficients. However, the choice of functionals has implications for the quality of the approximation as well as the computation required to obtain it. Combining equations (13), (14), and (15) we have

$$\phi_i\left(\mathbf{M}\sum_{j=1}^{n} \alpha_j u_j - g\right) = 0, \tag{16}$$

for $i = 1, 2, \ldots, n$, which is a system of n equations for the unknown coefficients $\alpha_1, \ldots, \alpha_n$. Exploiting the linearity of ϕ_i and $\mathbf{M}$, we can express the above equations in matrix form:

$$\begin{bmatrix} \phi_1\mathbf{M}u_1 & \cdots & \phi_1\mathbf{M}u_n \\ \vdots & \ddots & \vdots \\ \phi_n\mathbf{M}u_1 & \cdots & \phi_n\mathbf{M}u_n \end{bmatrix} \begin{bmatrix} \alpha_1 \\ \vdots \\ \alpha_n \end{bmatrix} = \begin{bmatrix} \phi_1 g \\ \vdots \\ \phi_n g \end{bmatrix}. \tag{17}$$

As shown below, most global illumination algorithms described in recent literature are special cases of the formulation in equation (17). Specific projection-based algorithms are characterized by the following choices:

1. Choice of basis functions $u_1, \ldots, u_n$
2. Choice of linear functionals $\phi_1, \ldots, \phi_n$
3. Algorithms for evaluating $\phi_i\mathbf{M}u_j$ for $i, j = 1, 2, \ldots, n$
4. Algorithms for solving the discrete linear system

These choices do not necessarily coincide with the sequential steps of an algorithm. Frequently information obtained during evaluation of the linear functionals or during the solution of the discrete linear system is used to alter the choice of basis functions. The essence of adaptive meshing lies in this form of feedback. Regardless of the order in which the steps are carried out, the approximation generated ultimately rests upon specific choices in each of the above categories. Consequently, we can determine conservative bounds on the accuracy of the final solution by studying the impact of these choices independently.

4 A Taxonomy of Errors

In the previous section we characterized the fundamental features that distinguish projection-based global illumination algorithms. In this section we introduce a higher-level organization motivated by distinct categories of error; this subsumes the previous ideas and adds the notion of imprecise problem instances. Assuming that accuracy is measured by comparing with the exact solution to equation (1), all sources of error incurred by projection methods fall into one of three categories:

- **Perturbed Boundary Data:**

  Both the operator $\boldsymbol{M}$ *and the source term* g *may be inexact due to measurement errors and/or simplifications made for efficiency.*

- **Discretization Error:**

  The finite-dimensional space $\mathcal{X}_n$ *may not include the exact solution. In addition, satisfying the constraints* $\phi_1, \ldots, \phi_n$ *may not select the best possible approximation from* $\mathcal{X}_n$.

- **Computational Errors:**

  The matrix elements $\phi_i \boldsymbol{M} u_j$ *may not be computed exactly, thus perturbing the discrete linear system. Finally, the perturbed linear system may not be solved exactly.*

It is important to note that the above categories are mutually exclusive and account for all types of errors incurred in solving equation (1) with a projection method. The conceptual error taxonomy is shown schematically in Figure 3. In the remainder of this section we illustrate each of these categories of error with examples from existing algorithms.

4.1 Perturbed Boundary Data

The idealized problems that we solve in practice are rarely as realistic as we would like. As a rule, we settle for solving "near by" problems for several reasons. First, the data used as input may only be approximate. For instance, reflectance and emission functions obtained through simulation [40] or empirically through measurement of actual materials or light sources [36, 39] are inherently contaminated by error.

A second reason that boundary data may be perturbed is that use of the exact data may be prohibitively expensive or even impossible. Thus, a surface may be treated as a smooth Lambertian reflector for the purpose of simulation, although the actual geometry and material exhibits directional diffuse scattering. Regardless of the source of the discrepancy, the near-by problem can be viewed as a *perturbation* of the original problem.

4.2 Discretization Error

To make the problem of global illumination amenable to computer solution, we must recast the problem in terms of finite-dimensional quantities and finite processes. This transition is referred to as *discretization*. In general, the discrete finite-dimensional problem cannot entirely capture the behavior of the infinite-dimensional problem, and the discrepancy is called *discretization error*. This type of error is particularly difficult to analyze as it inherently involves infinite-dimensional spaces and their relationship to finite-dimensional subspaces. Consequently, most global illumination algorithms rely upon heuristics rather than error bounds to perform discretization tasks such as adaptive meshing.

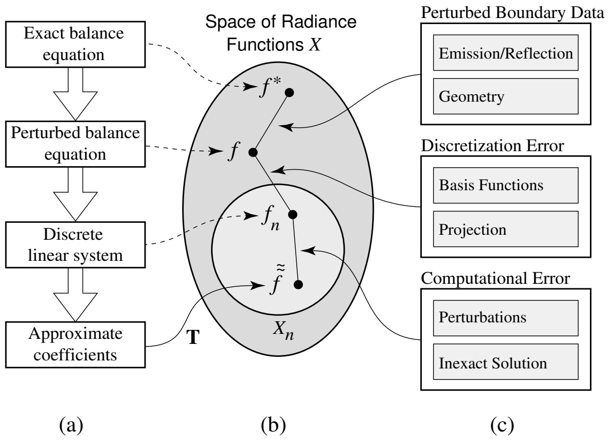

Figure 3: *(a) The conceptual stages for computing an approximate solution. (b) The function corresponding to each stage is an approximation of the previous stage. (c) Approximations specific to each stage introduce new errors.*

We now look at how two aspects of discretization have been treated in various algorithms: 1) the choice of basis functions, which determines the finite-dimensional space containing the approximation, and 2) projection, the finite process by which the approximation is selected from this space.

4.2.1 Basis Functions

For many global illumination algorithms, the basis functions are completely determined by the geometry of the boundary elements. This is true of any piecewise-constant approximation, such as those employed by Goral et al. [10], Cohen et al.[8], and Hanrahan et al. [13]. Often a smoothing step is applied to piecewise-constant approximations for display purposes, although this does not necessarily improve the accuracy of the approximation. When non-constant elements are employed, such as tensor product polynomials [41, 37] or wavelets [12], the degrees of freedom of the elements add to the dimension of the approximating subspace.

For non-diffuse environments the finite-dimensional space must also account for the directional dependence of surface radiance functions; several avenues have been explored for doing this. Immel et al. [17] used subdivided cubes centered at a finite number of surface points to simultaneously discretize directions and positions. Sillion et al. [35] used a truncated series of spherical harmonics to capture directional dependence and a quadrilateral mesh of surface elements for the spatial dependence. As a third contrasting approach, Aupperle et al. [3] used piecewise-constant functions defined over pairs of patches to account for both directional and spatial variations.

The accuracy of the approximation is limited by the space of basis functions; in general, the exact solution cannot be formed by a finite linear combination of polynomials or other basis functions. The error can be reduced either by expanding the subspace $\mathcal{X}_n$, or by selecting basis functions that fit the solution more closely. Discontinuity meshing [14, 24] is an example of the latter strategy.

4.2.2 Projections

Given a set of basis functions, the next conceptual step is to construct a reasonable approximation from it. We now demonstrate how the major projection methods follow from specific choices of the linear functionals $\phi_1, \ldots, \phi_n$ described in section 3. The first technique is *collocation* [22], which follows by defining ϕ_i to be an *evaluation functional* at the i^{th} collocation point; that is

$$\phi_i h \equiv h(x_i), \tag{18}$$

where $x_1, \ldots, x_n$ are distinct points in the domain of the radiance functions chosen so that $\det[u_i(x_j)] \neq 0$. The ij^{th} element of the resulting linear system has the form

$$u_j(x_i) - (\mathbf{KG}u_j)(x_i). \tag{19}$$

Collocation has been widely used in global illumination because of the relative simplicity of the above expression. In the case of constant basis functions over planar boundary elements, the resulting matrix has a unit diagonal with point-to-area form factors off the diagonal, which is a widely used radiosity formulation [8]. In general, methods based on a finite number of point-to-area interactions are collocation methods.

A second technique is the *Galerkin method*, which follows by defining ϕ_i to be an inner product functional with the i^{th} basis function; that is

$$\phi_i h \equiv \langle\, u_i \mid h \,\rangle\,, \tag{20}$$

where the inner product $\langle \cdot \mid \cdot \rangle$ denotes the integral of the product of two functions. The ij^{th} element of the resulting linear system has the form

$$\langle\, u_i \mid u_j \rangle - \langle\, u_i \mid \mathbf{KG}u_j \rangle\,. \tag{21}$$

The Galerkin method was first employed in global illumination by Goral et al. [10] using a uniform mesh and constant basis functions. The use of higher-order basis functions was investigated by Heckbert [14] and later by Zatz [41] and Troutman and Max [37]. For diffuse environments with constant elements, the second inner product in equation (21) reduces to an area-to-area form factor.

There are other possibilities for the linear functionals $\phi_1, \ldots, \phi_n$. For instance, the inner products in equation (20) may be taken with respect to a different set of basis functions. If we set

$$\phi_i h \equiv \langle\, \mathbf{M}u_i \mid h \,\rangle\,, \tag{22}$$

we obtain the *least squares* method. With the above functional, the solution to the linear system in equation (17) has a residual error that is orthogonal to the space $\mathcal{X}_n$, which minimizes the residual error in a very natural sense. However, the matrix elements for the least squares method include terms of the form $\langle\, \mathbf{KG}u_i \mid \mathbf{KG}u_j \rangle$, which are formidable to evaluate even with trivial basis functions [15]. Consequently, to our knowledge, there are currently no global illumination algorithms based on the least squares method.

4.3 Computational Errors

Given a particular method of discretization, error may be incurred in constructing the discrete (finite-dimensional) linear system, and once formulated, we may fail to solve even the discrete problem exactly. These facts illustrate a third distinct class of errors. Because they arise from the practical limits of computational procedures, this class is called *computational errors* [23]. Computational errors perturb the discrete problem, and then preclude exact solution of the perturbed problem. We now look at examples of each of these.

4.3.1 Perturbation of the Linear System

The most computationally expensive operation of global illumination is the evaluation of the matrix elements in equation (17); this is true even of algorithms that do not store an explicit matrix [7]. Furthermore, only in very special cases can the matrix elements be formed exactly, as they entail visibility calculations coupled with multiple integration. Consequently, the computed matrix is nearly always perturbed by computational errors.

A common example of an error in this category is imprecise form factors. The algorithm introduced by Cohen and Greenberg [8] for diffuse environments computed form factors at discrete surface points by means of a hemicube, which introduced a number of errors specific to this approach [5]. Errors are also introduced when form factors are approximated through ray tracing [38] or by using simpler geometries [13]. These errors can be mitigated to some extent by the use of analytic form factors [5, 34], yet there remain many cases for which no analytic expression is known, particularly in the presence of occluders.

For non-constant basis functions the form factor computations are replaced by more general inner products, which require different approximations. For instance, the matrix elements in the Galerkin approach of Zatz [41] required four-fold integrals, which were approximated using Gaussian quadrature. Non-diffuse environments pose a similar difficulty in that the matrix elements entail integration with reflectance functions [35].

Another form of matrix perturbation arises from simplifications made for the sake of efficiency. For example, small entries may be set to zero [25], or the entire matrix may be approximated by one with a more efficient structure, such as a block matrix [13] or a wavelet decomposition [12].

4.3.2 Inexact Solution of the Linear System

Once the discrete linear system is formed, we must solve for the coefficients $\alpha_1, \ldots, \alpha_n$. This has been done in a number of ways, including Gaussian elimination [10], Gauss-Seidel [8], Southwell relaxation (shooting) [7, 11], and Jacobi iteration [13]. In any such method, there will be some error introduced in the solution process: direct solvers like Gaussian elimination are prone to roundoff error, whereas interative solvers like Gauss-Seidel must halt after a finite number of iterations.

In approaches where the matrix is constructed in advance, such as "full matrix" radiosity [8] or hierarchical radiosity [13], the iterative solution can be carried out to essentially full convergence. Other approaches, such as progressive radiosity [7], construct matrix elements on the fly and then discard them. The cost of computing these elements generally precludes complete convergence, making this source of error significant.

5 Quantifying Error

Thus far the discussion has focused on the sources of error without attempting to quantify them. In this section we introduce the tools that will be used to derive error bounds in subsequent sections.

To quantify error, we require some notion of "size" for the elements of the function space $\mathcal{X}$ as well as a notion of the "distance" between its members. An abstraction that embodies these notions is a real-valued *norm* defined on $\mathcal{X}$, denoted by $\|\;\|$; the ordered pair $(\mathcal{X}, \|\;\|)$ is then called a *normed linear space*. The subject of normed linear spaces is extremely rich, and we shall only touch upon it here. For a more complete discussion see the classic texts by Rudin [32] or Kantorovich [19].

5.1 Function Norms

Infinitely many norms can be defined on the space of radiance functions, each conveying a different notion of size, distance, and convergence. Since equation (8) imposes only certain integrability conditions on radiance functions, it is natural to apply the L^p-norm defined by

$$\| f \|_p \equiv \left[\int_{\mathcal{M}} \int_{\mathcal{S}^2} |f(x, \omega)|^p \; d\mu(\omega)\, d\sigma(x) \right]^{1/p}, \qquad (23)$$

where σ denotes area measure, and p is a real number in the range $[1, \infty]$. Three of these norms are of particular interest: the L^1-norm and L^∞-norm have immediate physical interpretations, while the L^2-norm possesses algebraic properties that make it appropriate in some instances; for instance, when the algebraic structure of an inner product space is required.

In the limiting case of $p = \infty$, the L^p-norm reduces to

$$\| f \|_\infty \equiv \sup_{x \in \mathcal{M}} \sup_{\omega \in \mathcal{S}^2} |f(x, \omega)|, \qquad (24)$$

where sup denotes the least upper bound. As the *maximum radiance* attained at any surface point and in any direction, $\| f \|_\infty$ has the dimensions of radiance [watts/m^{2}sr]. In contrast, $\| f \|_1$ measures the *total power* of the radiance function f, and consequently has the dimensions of power [watts].

To measure the distance between two radiance functions f and f' in $\mathcal{X}$, we use the non-negative quantity $\| f - f' \|$. Similarly, we define the distance between a function and a subspace $Y \subset \mathcal{X}$ by

$$\mathrm{dist}(f, Y) \equiv \inf_{f' \in Y} \left\| f - f' \right\|, \qquad (25)$$

where inf denotes the greatest lower bound. When the subscript on the norm is omitted, it is implied that the definition or relation holds for any choice of norm.

5.2 Operator Norms

To investigate the effects of perturbing the operators themselves, we must also endow the linear space of operators on $\mathcal{X}$ with a norm. The standard *operator norm* is defined by

$$\| \mathbf{A} \| \equiv \sup \{ \| \mathbf{A}h \| \, : \, \| h \| \leq 1 \}, \qquad (26)$$

where the norm appearing twice on the right is the one associated with $\mathcal{X}$; the operator norm is said to be *induced* by the given function norm. Although the theory of linear operators closely parallels that of matrices, there are important differences; for instance, matrix norms are necessarily finite while operator norms need not be. We therefore distinguish the class of *bounded* operators as those with finite norm.

Equation (26) implies that $\| \mathbf{A}h \| \leq \| \mathbf{A} \| \, \| h \|$ for all $h \in \mathcal{X}$. Several useful consequences follow immediately from this inequality. For instance, let $f_n \in \mathcal{X}_n$ be the approximate solution to $\mathbf{M}f = g$ discussed in section 3. Then, whenever $\mathbf{M}^{-1}$ exists and is bounded, it follows from equation (14) that

$$\begin{aligned} \| f_n - f \| &= \left\| f_n - \mathbf{M}^{-1} g \right\| \\ &\leq \left\| \mathbf{M}^{-1} \right\| \| \mathbf{M} f_n - g \| \\ &\leq \left\| \mathbf{M}^{-1} \right\| \| r_n \|, \end{aligned} \qquad (27)$$

which justifies the strategy of minimizing the residual error r_n in order to reduce the error $\| f_n - f \|$. Appendix A contains several other inequalities, including Banach's lemma, which will prove useful in deriving error bounds.

5.3 Norms of Special Operators K, G, and $\mathbf{M}^{-1}$

We now compute bounds for the operators $\mathbf{K}$, $\mathbf{G}$, and $\mathbf{M}^{-1}$, which will be essential for all subsequent bounds. From definitions (5), (23), and (26), we can deduce an explicit formula for $\| \mathbf{K} \|_1$. A straightforward manipulation [19, p. 109] shows that

$$\| \mathbf{K} \|_1 = \sup_{x \in \mathcal{M}} \sup_{\omega' \in \Omega_i} \int_{\Omega_o} k(x; \omega' \to \omega)\, d\mu(\omega). \qquad (28)$$

This norm is the maximum reflectivity of any surface in the environment which, by equation (3), is necessarily bounded above by 1. If we disallow perfect reflectors and ignore the wave optics effect of specular reflection near grazing, the norm has a bound strictly less than 1; that is, $\| \mathbf{K} \|_1 = m < 1$. The L^∞-norm of $\mathbf{K}$ is very similar, merely exchanging the roles of the two directions:

$$\| \mathbf{K} \|_\infty = \sup_{x \in \mathcal{M}} \sup_{\omega \in \Omega_o} \int_{\Omega_i} k(x; \omega' \to \omega)\, d\mu(\omega'). \qquad (29)$$

By the reciprocity relation in equation (4) we have $\| \mathbf{K} \|_1 = \| \mathbf{K} \|_\infty$, implying that reflected radiance is everywhere diminished by at least the same factor of m. The bound can be extended to all L^p-norms by the inequality

$$\| \mathbf{A} \|_p \leq \max \left\{ \| \mathbf{A} \|_1, \| \mathbf{A} \|_\infty \right\}, \qquad (30)$$

which holds for linear integral operators [20, p. 144]. Therefore

$$\| \mathbf{K} \|_p \leq m < 1 \qquad (31)$$

for all $1 \leq p \leq \infty$.

Similarly, we can bound the $\mathbf{G}$ operator by invoking several physical principles. Because we have assumed no volume emission or absorption, the enclosure $\mathcal{M}$ contains no sources or sinks in its interior. By Gauss's theorem, an analogous statement holds for the boundary of the enclosure; that is, the net flux passing through $\mathcal{M}$ is zero. Because the power entering the enclosure must equal the power leaving it, we have $\| f \|_1 = \| \mathbf{G}f \|_1$. The argument extends immediately to f^p for any p, so $\| f \|_p = \| \mathbf{G}f \|_p$. By definition of the operator norm, it follows that

$$\| \mathbf{G} \|_p = 1 \qquad (32)$$

for all $1 \leq p \leq \infty$. Note that the norm can be less than 1 if $\mathcal{M}$ does not form an enclosure.

A third operator that we shall need to bound is $\mathbf{M}^{-1}$, which maps surface emission functions to surface radiance functions at equilibrium. It follows from bounds (31) and (32) that $\mathbf{M}^{-1}$ exists and can be expanded into a Neumann series [20, p. 30]. Taking norms and summing the resulting geometric series, we have

$$\left\| \mathbf{M}^{-1} \right\|_p \leq 1 + m + m^2 + \ldots = \frac{1}{1 - m} \qquad (33)$$

for all $1 \leq p \leq \infty$. This bound also implies that the space of bounded radiance functions with respect to a given L^p-norm is closed under global illumination; that is, if an emission function g has a finite L^p-norm, so too will the solution f.

6 Error Bounds

With the results of the previous section we can obtain bounds for each category of error. Because the methods employed here take account of very little information about an environment, the bounds tend to be quite conservative. The organizational structure of this section follows that of Figure 3 and section 4.

6.1 Perturbed Boundary Data

We solve global illumination problems using inexact or noisy data in hopes of obtaining a solution that is close to that of the original problem. But under what circumstances is this a reasonable expectation? To answer this question we analyze the mapping from problem instances $(\mathcal{M}, \mathbf{M}, g)$ to solutions f. We shall assume that any input to a global illumination problem may be contaminated by error, and determine the impact of these errors on the solution. We shall show that the problem of global illumination is *well-posed* except in certain extreme and non-realizable cases; that is, "small" perturbations of the input data produce "small" errors in the solution in physically meaningful problems.

To bound the effects of input data perturbations, we examine the quantity $\| f^* - f \|$, where f^* is the solution to the exact or unperturbed system, and f is the solution to a perturbed system

$$\widetilde{\mathbf{M}} f = \widetilde{g}, \tag{34}$$

where perturbed entities are denoted with tildes.

6.1.1 Perturbed Reflectance and Emission Functions

We first investigate the effect of perturbing the reflectance and emission functions. The former is equivalent to perturbing the local reflection operator $\mathbf{K}$. Consider a perturbed operator $\widetilde{\mathbf{K}}$ and a perturbed emission function $\widetilde{g}$ such that

$$\left\| \mathbf{K} - \widetilde{\mathbf{K}} \right\| \leq \delta_k \tag{35}$$

and

$$\| g - \widetilde{g} \| \leq \delta_g. \tag{36}$$

Since $\| \mathbf{G} \| = 1$, and $\mathbf{G}$ is not perturbed, it follows that inequality (35) also applies to the corresponding perturbed $\mathbf{M}$ operator:

$$\left\| \widetilde{\mathbf{M}} - \mathbf{M} \right\| = \left\| \mathbf{KG} - \widetilde{\mathbf{K}}\mathbf{G} \right\| \leq \left\| \mathbf{K} - \widetilde{\mathbf{K}} \right\| \, \| \mathbf{G} \| \leq \delta_k.$$

Intuitively, the above inequality holds because the worst-case behavior over all possible field radiance functions is unaffected by $\mathbf{G}$, which merely redistributes radiance.

A bound on $\| f^* - f \|$ in terms of bounds on the perturbations and other quantities is derived in Appendix B. The result is

$$\| f^* - f \| \leq \left(\frac{\delta_k}{1 - m - \delta_k} \right) \left(\frac{\| g \| + \delta_g}{1 - m} \right) + \frac{\delta_g}{1 - m}, \tag{37}$$

which contains terms accounting for perturbations in the reflectance and emission functions individually as well as a second-order term involving $\delta_k \delta_g$. Note that the reflectance term requires $\delta_k < 1 - m$, indicating that the problem becomes less stable as the maximum reflectivity approaches 1. For highly reflective environments, the results may be arbitrarily bad if the input data are not correspondingly accurate. With perfect reflection the problem is ill-posed. In general, the worst-case absolute error in f depends upon the maximum reflectivity of the environment m, the perturbation of the reflection functions δ_k, and the error in the emission function, δ_g.

6.1.2 Perturbed Surface Geometry

The effects of imprecise surface geometry are more difficult to analyze than those due to imprecise reflection or emission. While the space of radiance functions has a linear algebraic structure that underlies all of the analysis, no analogous structure exists on the set of possible surface geometries. The analysis must therefore proceed along different lines.

An alternative approach is to study imprecise surface geometry indirectly, by means of the $\mathbf{G}$ operator. That is, we can express the effect of surface perturbations on the field radiance at each point as a perturbation of $\mathbf{G}$. Given a perturbed operator $\widetilde{\mathbf{G}}$, and a bound on its distance from the exact $\mathbf{G}$, the same analysis used in Appendix B for perturbed reflectance functions can be applied. Such an approach may be useful for analyzing schemes in which geometry is simplified to improve the efficiency of global ilumination [33]. However, relating changes in geometry to bounds on the perturbation of $\mathbf{G}$ is an open problem.

6.2 Discretization Error

In this section we study discretization errors introduced by projection methods. Clearly, the discretization error $\| f - f_n \|$ is bounded from *below* by $\mathrm{dist}(f, \mathcal{X}_n)$, the distance to the best approximation attainable within the space $\mathcal{X}_n$. To obtain an upper bound, we express equation (15) using an explicit projection operator:

$$\mathbf{P}_n r_n = 0, \tag{38}$$

where $\mathbf{P}_n$ projects onto the subspace $\mathcal{X}_n$. That is, $\mathbf{P}_n$ is a linear operator with $\mathbf{P}_n^2 = \mathbf{P}_n$ and $\mathbf{P}_n h = h$ for all $h \in \mathcal{X}_n$. Such a projection can be defined in terms of the basis functions $u_1, \ldots, u_n$ and the linear functionals $\phi_1, \ldots, \phi_n$ of section 3. The form of $\mathbf{P}_n$ is particularly simple when

$$\phi_i(u_j) = \begin{cases} 1 & \text{if } i = j \\ 0 & \text{otherwise} \end{cases} \tag{39}$$

which is commonly the case. For instance, this holds whenever there is exactly one collocation point within the support of each basis function, or when orthogonal polynomials are used in a Galerkin-based approach. When equation (39) holds, $\mathbf{P}_n$ is given by

$$\mathbf{P}_n h = \sum_{i=1}^{n} \phi_i(h) u_i \tag{40}$$

for any function $h \in \mathcal{X}$ [9]. It is easy to see that the above equation defines a projection onto $\mathcal{X}_n$.

To produce the desired bound, we write equation (38) as

$$\mathbf{P}_n \mathbf{M}(f - f_n) = 0, \tag{41}$$

then isolate the quantity $\| f - f_n \|$ after taking the norm, as shown in Appendix C. The resulting bounds on discretization error are

$$\mathrm{dist}(f, \mathcal{X}_n) \leq \| f - f_n \| \leq \left(\frac{\mathrm{dist}(f, \mathcal{X}_n)}{1 - m - \delta_P} \right) \left(1 + \| \mathbf{P}_n \| \right), \tag{42}$$

where the constant δ_P is such that

$$\| \mathbf{K} - \mathbf{P}_n \mathbf{K} \| \leq \delta_P. \tag{43}$$

The bounds in equation (42) depend on both the subspace $\mathcal{X}_n$ and the projection method. Note that if f is in the space $\mathcal{X}_n$, then $\mathrm{dist}(f, \mathcal{X}_n) = 0$, and the upper bound implies that $f_n = f$. Thus, all projection methods find the exact solution when it is achievable with a linear combination of the given basis functions. On the other hand, when $\mathrm{dist}(f, \mathcal{X}_n)$ is large, then the lower bound implies that the approximation will be poor even when all other steps are exact. Unfortunately, this distance is difficult to estimate *a priori*, as it depends on the actual solution.

The dependence on the type of projection appears in the factor of $1 + \| \mathbf{P}_n \|$ and in the constant δ_P. For all projections based on inner products, $\| \mathbf{P}_n \| = 1$ when the basis functions are orthogonal [2, p. 64]. For other methods, such as collocation, the norm of the projection may be greater than one [2, p. 56]. The meaning of the constant δ_P is more subtle. The norm in equation (43) is a measure of how well the projection $\mathbf{P}_n$ captures features of the reflected radiance.

6.3 Computational Errors

The effects of computational errors can be estimated by treating them as perturbations of the *discrete* linear system. The analysis therefore parallels that of perturbed boundary data, although carried out in a finite dimensional space. We shall denote the linear system in equation (17) by $\mathbf{A}\alpha = \mathbf{b}$. In general, the matrix elements as well as the vector $\mathbf{b}$ will be inexact due to errors or simplifications. We denote the perturbed system and its solution by

$$\widetilde{\mathbf{A}}\,\widetilde{\alpha} = \widetilde{\mathbf{b}}. \tag{44}$$

Although the exact matrix is unknown, it is frequently possible to bound the error present in each element. For instance, this can be done for approximate form factors and block matrix approximations [13].

Given element-by-element error bounds, the impact on the final solution can be bounded. From the element perturbations we can find δ_A and δ_b such that

$$\left\| \, \mathbf{A} - \widetilde{\mathbf{A}} \, \right\| \leq \delta_A, \tag{45}$$

and

$$\left\| \, \mathbf{b} - \widetilde{\mathbf{b}} \, \right\| \leq \delta_b, \tag{46}$$

for some vector norm and induced matrix norm. Also, since the perturbed matrix is known, the norm of its inverse can be estimated:

$$\left\| \, \widetilde{\mathbf{A}}^{-1} \, \right\| \leq \beta. \tag{47}$$

With the three bounds described above, essentially the same steps used in Appendix B yield the bound

$$\| \, \alpha - \widetilde{\alpha} \, \| \leq \left(\frac{\delta_A \beta^2}{1 - \delta_A \beta} \right) \left(\left\| \, \widetilde{\mathbf{b}} \, \right\| + \delta_b \right) + \beta \delta_b. \tag{48}$$

Computing values for δ_A and δ_b that are reasonably tight is almost always difficult, requiring error bounds for each step in forming the matrix elements. There is no universal method by which this can be done; each approach to estimating visibility or computing inner products, for example, requires a specialized analysis. In contrast, the bound in equation (47) is more accessible, as it is purely a problem of linear algebra.

Given that equation (44) generally cannot be solved exactly, the solution process is yet another source of error. The result is an approximation of $\widetilde{\alpha}$, which we denote by $\widetilde{\widetilde{\alpha}}$. From equation (27), this last source of error is bounded by

$$\left\| \, \widetilde{\alpha} - \widetilde{\widetilde{\alpha}} \, \right\| \leq \beta \left\| \, \widetilde{\mathbf{A}} \widetilde{\widetilde{\alpha}} - \widetilde{\mathbf{b}} \, \right\|. \tag{49}$$

When $\widetilde{\mathbf{A}}$ and $\widetilde{\mathbf{b}}$ are stored explicitly, the above expression can be used as the stopping criterion for an iterative solver.

To relate errors in the coefficients $\alpha_1, \ldots, \alpha_n$ to errors in the resulting radiance function, consider the mapping $\mathbf{T} : \mathbb{R}^n \rightarrow \mathcal{X}_n$ where

$$\mathbf{T}x \equiv \sum_{j=1}^{n} x_j \, u_j. \tag{50}$$

As a finite-dimensional linear operator, $\mathbf{T}$ is necessarily bounded; its norm supplies the connection between coefficients in $\mathbb{R}^n$ and functions in $\mathcal{X}_n$. Observe that

$$\begin{aligned} \left\| \, f_n - \widetilde{\widetilde{f}} \, \right\| &= \left\| \, \mathbf{T}\alpha - \mathbf{T}\widetilde{\widetilde{\alpha}} \, \right\| \\ &\leq \| \mathbf{T} \| \left\| \, \alpha - \widetilde{\widetilde{\alpha}} \, \right\| \\ &\leq \| \mathbf{T} \| \left(\left\| \, \alpha - \widetilde{\alpha} \, \right\| + \left\| \, \widetilde{\alpha} - \widetilde{\widetilde{\alpha}} \, \right\| \right). \end{aligned} \tag{51}$$

Equation (51) relates the computational error present in the final solution to inequalities (48) and (49). The value of $\| \mathbf{T} \|$ will depend on the basis functions $u_1, \ldots, u_n$ and the choice of norms for both $\mathbb{R}^n$ and $\mathcal{X}_n$, which need not be related.

7 The Combined Effect of Errors

In the previous sections we derived inequalities to bound the errors introduced into the solution of a global illumination problem. Using these inequalities we can now bound the distance between the exact solution and the computed solution. By the triangle inequality we have

$$\left\| \, f^* - \widetilde{\widetilde{f}} \, \right\| \leq \left\| \, f^* - f \, \right\| + \left\| \, f - f_n \, \right\| + \left\| \, f_n - \widetilde{\widetilde{f}} \, \right\|, \tag{52}$$

which is the numerical analogue of the chain of approximations shown in Figure 3. The terms on the right correspond to errors arising from perturbed boundary data, discretization, and computation; sections 6.1, 6.2, and 6.3 provide bounds for each of these errors. We note that the first and third terms can be reduced in magnitude by decreasing errors in the emission and reflectance functions and in the computational methods for forming and solving the linear system. The second term, due to discretization error and addressed by equation (42), is more difficult to compute, and is further explored below.

7.1 A Numerical Experiment

Using a simple configuration, we compared the discretization errors incurred by two projection methods: collocation and Galerkin. The geometry consisted of a square diffuse light source, with unit side, above and parallel to a rectangular diffuse receiver 10 units on a side. The receiver was discretized using piecewise-constant basis functions over uniform quadrilateral elements. The reflectivity of the light source was set to zero, so interreflection was ignored. For this configuration the projections and the actual solution could be computed analytically using closed-form expressions for point-to-area and area-to-area form factors [16].

Figure 4 shows various errors with respect to three different norms as the number of mesh elements increases. The bottom curve in each plot is the distance between the closest approximation and the analytic solution; by equation (42) this is a lower bound on the the discretization error of any projection method. For each mesh element, the closest approximation follows from the median, the mean, and the extrema of the analytic solution for the L^1, L^2, and L^∞-norms, respectively. Also shown in the plots are the distances from the analytic solution to the approximations produced by the two projections.

Because the constants m and δ_P in equation (42) are related to reflectivity, which did not enter into this computation, the theoretical upper bound on discretization error is the lower bound times $1 + \| \mathbf{P} \|$. The top curve in each plot is twice the lower bound, which is the smallest upper bound attainable by the present analysis since $\| \mathbf{P} \| \geq 1$ for any projection. With one exception, the top curve suffices as an upper bound for the discretization error. (The single exception results from a collocation projection with a norm significantly greater than 1.)

The discretization error is one of the more difficult errors to quantify. Note that the discretization errors of the two projections shown in Figure 4 generally fall between the upper and lower curves, and approach the lower bound as the mesh refinement is increased. The two projection methods are very close, and the Galerkin projection exactly achieves the lower bound in the L^2-norm.

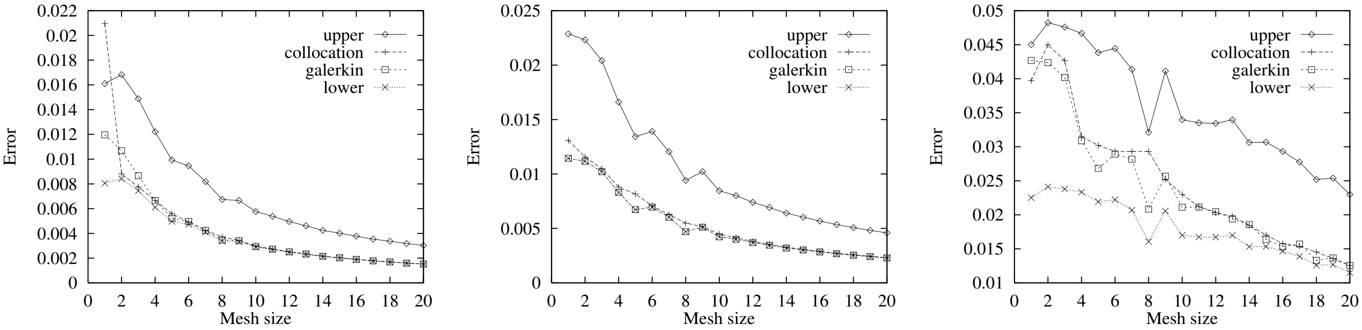

Figure 4: *Discretization errors in three different norms for collocation and Galerkin applied to a simple configuration: (a) L^1-norm, (b) L^2-norm, and (c) L^∞-norm. The lower curve in each plot is the distance to the optimal approximation, and the upper curve is twice this distance.*

8 Conclusions and Future Work

We have introduced a new operator equation describing direct radiative transfer among opaque surfaces. The new formulation is appropriate for global illumination, and has several advantages over the rendering equation. First, it is closer to the underlying physical principles, being based on standard radiometric concepts. Second, it is well-suited to the error analysis presented in this paper.

We have identified three sources of error in global illumination algorithms: inaccuracies in the input data, discretization, and computational errors. Input errors result from noise in measurement or simulation, or from simplifications. Discretization errors result from restricting the space of possible approximations and from the method of selecting the approximation. Computational errors form a large class that includes imprecise form factors and visibility, as well as errors introduced by block matrix approximations and iterative matrix methods. To produce a reliable solution, each of these sources of error must be accounted for. Using standard methods of analysis, we have derived worst-case bounds for each category of error based primarily on quantities such as the maximum directional reflectivity of any surface in the environment.

There remains much work to do. Better ways of characterizing environments are needed to obtain more realistic error bounds; constants such as maximum reflectivity are much too coarse. A deficiency of the function norms employed here is that they do not adequately handle the wave optics effect of specular reflection at grazing angles. Other norms should be explored, including those that are in some sense perceptually-based. Also, reliable bounds are needed for a wide assortment of standard computations, such as form factors between partially occluded surfaces and inner products involving higher-order elements. Finally, we need to relate the norms used here and the types of tolerances that can reasonably be obtained through measurement and simulation.

Acknowledgments

The authors wish to thank Dani Lischinski for many valuable discussions. This work was supported by the NSF grant "Interactive Computer Graphics Input and Display Techniques" (CCR-8617880), and by the NSF/ARPA Science and Technology Center for Computer Graphics and Scientific Visualization (ASC-8920219). The authors gratefully acknowledge the generous equipment grant from Hewlett-Packard Corporation on whose workstations this research was conducted.

REFERENCES

[1] ANSELONE, P. M. Convergence and error bounds for approximate solutions of integral and operator equations. In *Error in Digital Computation*, L. B. Rall, Ed., vol. 2. John Wiley & Sons, 1965, pp. 231–252.

[2] ATKINSON, K. E. *A Survey of Numerical Methods for the Solution of Fredholm Integral Equations of the Second Kind*. Society for Industrial and Applied Mathematics, Philadelphia, 1976.

[3] AUPPERLE, L., AND HANRAHAN, P. A hierarchical illumination algorithm for surfaces with glossy reflection. In *Computer Graphics Proceedings* (1993), Annual Conference Series, ACM SIGGRAPH, pp. 155–162.

[4] BATEMAN, H. Report on the history and present state of the theory of integral equations. *Report of the 18th Meeting of the British Association for the Advancement of Science* (1910), 345–424.

[5] BAUM, D. R., RUSHMEIER, H. E., AND WINGET, J. M. Improving radiosity solutions through the use of analytically determined form-factors. *Computer Graphics 23*, 3 (July 1989), 325–334.

[6] CHANDRASEKAR, S. *Radiative Transfer*. Dover Publications, New York, 1960.

[7] COHEN, M. F., CHEN, S. E., WALLACE, J. R., AND GREENBERG, D. P. A progressive refinement approach to fast radiosity image generation. *Computer Graphics 22*, 4 (August 1988), 75–84.

[8] COHEN, M. F., AND GREENBERG, D. P. The hemi-cube: A radiosity solution for complex environments. *Computer Graphics 19*, 3 (July 1985), 75–84.

[9] GOLBERG, M. A. A survey of numerical methods for integral equations. In *Solution methods for integral equations: Theory and applications*, M. A. Golberg, Ed. Plenum Press, New York, 1979, pp. 1–58.

[10] GORAL, C. M., TORRANCE, K. E., GREENBERG, D. P., AND BATTAILE, B. Modeling the interaction of light between diffuse surfaces. *Computer Graphics 18*, 3 (July 1984), 213–222.

[11] GORTLER, S. J., AND COHEN, M. F. Radiosity and relaxation methods. Tech. Rep. TR 408-93, Princeton University, 1993.

[12] GORTLER, S. J., SCHRÖDER, P., COHEN, M. F., AND HANRAHAN, P. Wavelet radiosity. In *Computer Graphics Proceedings* (1993), Annual Conference Series, ACM SIGGRAPH, pp. 221–230.

[13] HANRAHAN, P., SALZMAN, D., AND AUPPERLE, L. A rapid hierarchical radiosity algorithm. *Computer Graphics 25*, 4 (July 1991), 197–206.

[14] HECKBERT, P. S. *Simulating Global Illumination Using Adaptive Meshing*. PhD thesis, University of California, Berkeley, June 1991.

[15] HILDEBRAND, F. B. *Methods of Applied Mathematics*. Prentice-Hall, New York, 1952.

[16] HOWELL, J. R. *A Catalog of Radiation Configuration Factors*. McGraw-Hill, New York, 1982.

[17] IMMEL, D. S., COHEN, M. F., AND GREENBERG, D. P. A radiosity method for non-diffuse environments. *Computer Graphics 20*, 4 (August 1986), 133–142.

[18] KAJIYA, J. T. The rendering equation. *Computer Graphics 20*, 4 (August 1986), 143–150.

[19] KANTOROVICH, L., AND AKILOV, G. P. *Functional Analysis in Normed Spaces*. Pergamon Press, New York, 1964.

[20] KATO, T. *Perturbation Theory for Linear Operators*. Springer-Verlag, New York, 1966.

[21] KRASNOSEL'SKII, M. A., VAINIKKO, G. M., ZABREIKO, P. P., RUTITSKII, Y. B., AND STETSENKO, V. Y. *Approximate Solution of Operator Equations*. Wolters-Noordhoff, Groningen, The Netherlands, 1972.

Translated by D. Louvish.
[22] KRESS, R. *Linear Integral Equations*. Springer-Verlag, New York, 1989.
[23] LINZ, P. *Theoretical Numerical Analysis, an Introduction to Advanced Techniques*. John Wiley & Sons, New York, 1979.
[24] LISCHINSKI, D., TAMPIERI, F., AND GREENBERG, D. P. Discontinuity meshing for accurate radiosity. *IEEE Computer Graphics and Applications 12*, 6 (November 1992), 25–39.
[25] LISCHINSKI, D., TAMPIERI, F., AND GREENBERG, D. P. Combining hierarchical radiosity and discontinuity meshing. In *Computer Graphics Proceedings* (1993), Annual Conference Series, ACM SIGGRAPH, pp. 199–208.
[26] MACKERLE, J., AND BREBBIA, C. A., Eds. *The Boundary Element Reference Book*. Springer-Verlag, New York, 1988.
[27] MODEST, M. F. *Radiative Heat Transfer*. McGraw-Hill, New York, 1993.
[28] ORTEGA, J. M. *Numerical Analysis, a Second Course*. Academic Press, New York, 1972.
[29] PHILLIPS, J. L. The use of collocation as a projection method for solving linear operator equations. *SIAM Journal on Numerical Analysis 9*, 1 (1972), 14–28.
[30] PLANCK, M. *The Theory of Heat Radiation*. Dover Publications, New York, 1988.
[31] POLYAK, G. L. Radiative transfer between surfaces of arbitrary spatial distribution of reflection. In *Convective and Radiative Heat Transfer*. Publishing House of the Academy of Sciences of the USSR, Moscow, 1960.
[32] RUDIN, W. *Functional Analysis*. McGraw-Hill, New York, 1973.
[33] RUSHMEIER, H. E., PATTERSON, C., AND VEERASAMY, A. Geometric simplification for indirect illumination calculations. *Graphics Interface '93* (May 1993), 227–236.
[34] SCHRÖDER, P., AND HANRAHAN, P. On the form factor between two polygons. In *Computer Graphics Proceedings* (1993), Annual Conference Series, ACM SIGGRAPH, pp. 163–164.
[35] SILLION, F., ARVO, J., WESTIN, S., AND GREENBERG, D. P. A global illumination solution for general reflectance distributions. *Computer Graphics 25*, 4 (July 1991), 187–196.
[36] TOULOUKIAN, Y. S., Ed. *Retrieval Guide to Thermophysical Properties Research Literature*, second ed. McGraw-Hill, New York, 1968.
[37] TROUTMAN, R., AND MAX, N. L. Radiosity algorithms using higher-order finite element methods. In *Computer Graphics Proceedings* (1993), Annual Conference Series, ACM SIGGRAPH, pp. 209–212.
[38] WALLACE, J., ELMQUIST, K., AND HAINES, E. A ray tracing algorithm for progressive radiosity. *Computer Graphics 23*, 3 (July 1989), 315–324.
[39] WARD, G. J. Measuring and modeling anisotropic reflection. *Computer Graphics 26*, 2 (July 1992), 265–272.
[40] WESTIN, S., ARVO, J., AND TORRANCE, K. Predicting reflectance functions from complex surfaces. *Computer Graphics 26*, 2 (July 1992), 255–264.
[41] ZATZ, H. Galerkin radiosity: A higher order solution method for global illumination. In *Computer Graphics Proceedings* (1993), Annual Conference Series, ACM SIGGRAPH, pp. 213–220.

Appendix A: Operator-Norm Inequalities

In addition to the properties common to all norms, operator norms also satisfy

$$\| \mathbf{AB} \| \leq \| \mathbf{A} \| \; \| \mathbf{B} \| , \tag{53}$$

which makes them compatible with the multiplicative structure of operators [19]. Additional bounds, such as those pertaining to inverse operators, can be deduced from the basic properties of operator norms. For instance, given a bounded operator $\mathbf{A}$ with an inverse, any operator $\mathbf{B}$ sufficiently close to $\mathbf{A}$ is also invertible, with

$$\| \mathbf{B}^{-1} \| \leq \frac{\| \mathbf{A}^{-1} \|}{1 - \| \mathbf{A} - \mathbf{B} \| \, \| \mathbf{A}^{-1} \|}. \tag{54}$$

This inequality, known as Banach's lemma [28, p. 32], holds whenever $\| \mathbf{A} - \mathbf{B} \| < 1/\| \mathbf{A}^{-1} \|$. A useful corollary of Banach's lemma is the inequality

$$\| \mathbf{A}^{-1} - \mathbf{B}^{-1} \| \leq \frac{\| \mathbf{A} - \mathbf{B} \| \, \| \mathbf{A}^{-1} \|^2}{1 - \| \mathbf{A} - \mathbf{B} \| \, \| \mathbf{A}^{-1} \|}, \tag{55}$$

which holds under the same conditions [20, p. 31].

Appendix B: Error Bounds for Perturbed Reflectance and Emission Functions

To bound the error $\| f^* - f \|$ due to perturbations in the reflection and emission functions according to the inequalities (35) and (36) we write

$$\begin{aligned} \| f^* - f \| &= \left\| \mathbf{M}^{-1} g - \widetilde{\mathbf{M}}^{-1} \widetilde{g} \right\| \\ &\leq \left\| \mathbf{M}^{-1} - \widetilde{\mathbf{M}}^{-1} \right\| \; \| \widetilde{g} \| + \left\| \mathbf{M}^{-1} \right\| \; \| g - \widetilde{g} \| . \end{aligned}$$

From inequalities inequality (55), (33) and (35), we have

$$\left\| \mathbf{M}^{-1} - \widetilde{\mathbf{M}}^{-1} \right\| \leq \left(\frac{\delta_k}{1 - m - \delta_k} \right) \left(\frac{1}{1 - m} \right). \tag{56}$$

Combining the above and noting that $\| \widetilde{g} \| \leq \| g \| + \delta_g$ by inequality (36), we arrive at the bound

$$\| f^* - f \| \leq \left(\frac{\delta_k}{1 - m - \delta_k} \right) \left(\frac{\| g \| + \delta_g}{1 - m} \right) + \frac{\delta_g}{1 - m}. \tag{57}$$

Appendix C: Bounding Discretization Error

Adding $(\mathbf{I} - \mathbf{P}_n)(f - f_n)$ to both sides of equation (41) and simplifying, using the fact that $(\mathbf{I} - \mathbf{P}_n) f_n = 0$, we have

$$[\mathbf{I} - \mathbf{P}_n \mathbf{KG}] \, (f - f_n) = (\mathbf{I} - \mathbf{P}_n) f. \tag{58}$$

When the operator on the left of equation (58) is invertible, we obtain the bound

$$\| f - f_n \| \leq \left\| (\mathbf{I} - \mathbf{P}_n \mathbf{KG})^{-1} \right\| \; \| f - \mathbf{P}_n f \| . \tag{59}$$

A more meaningful bound is obtained by simplifying both factors on the righthand side [29]. Since $\mathbf{I} - \mathbf{P}_n \mathbf{KG}$ is an approximation of the operator $\mathbf{M}$, let δ_P be such that

$$\| \mathbf{M} - (\mathbf{I} - \mathbf{P}_n \mathbf{KG}) \| \leq \delta_P, \tag{60}$$

which simplifies to

$$\| \mathbf{K} - \mathbf{P}_n \mathbf{K} \| \leq \delta_P. \tag{61}$$

Because $\mathbf{M}$ is invertible, so is $\mathbf{I} - \mathbf{P}_n \mathbf{KG}$ when δ_P is sufficiently small. Banach's lemma from Appendix A then provides the bound

$$\left\| (\mathbf{I} - \mathbf{P}_n \mathbf{KG})^{-1} \right\| \leq \frac{1}{1 - m - \delta_P}.$$

The second norm on the right of inequality (59) can be simplified as follows. Let $h \in \mathcal{X}_n$. Then $\mathbf{P}_n h = h$, so

$$\begin{aligned} \| f - \mathbf{P}_n f \| &= \| (f - h) + (h - \mathbf{P}_n f) \| \\ &\leq \| f - h \| + \| \mathbf{P}_n (h - f) \| \\ &\leq (1 + \| \mathbf{P}_n \|) \, \| f - h \| . \end{aligned}$$

Since $h \in \mathcal{X}_n$ was chosen arbitrarily, the inequality holds for the least upper bound over $\mathcal{X}_n$, giving

$$\| f - \mathbf{P}_n f \| \leq (1 + \| \mathbf{P}_n \|) \, \mathrm{dist}(f, \mathcal{X}_n).$$

From the above inequalities we obtain the upper bound

$$\| f - f_n \| \leq \left(\frac{1}{1 - m - \delta_P} \right) \left(1 + \| \mathbf{P}_n \| \right) \, \mathrm{dist}(f, \mathcal{X}_n). \tag{62}$$