

Structured Penumbra Irradiance Computation

George Drettakis, Eugene Fiume

Abstract— A definitive understanding of irradiance behavior in penumbral regions has been hard to come by, mainly due to the computational expense of determining the visible parts of an area light source. Consequently, sampling strategies have been mostly ad hoc, and evaluation of the resulting approximations has been difficult. In this paper, the structure of penumbral irradiance is investigated empirically and numerically. This study has been made feasible by the use of the discontinuity mesh and the backprojection, an efficient data structure representing visibility in regions of partial occlusion. Regions of penumbræ in which irradiance varies non-monotonically are characterized empirically, and numerical tests are performed to determine the frequency of their occurrence. This study inspired the development of two algorithms for the construction of interpolating approximations to irradiance: one algorithm reduces the number of edges in the mesh defining the interpolant domain, and the other algorithm chooses among linear, quadratic, and mixed interpolants based on irradiance monotonicity. Results from numerical tests and images are presented that demonstrate good performance of the new algorithms for various realistic test configurations.

I. IRRADIANCE PROPERTIES IN SCENES WITH PARTIAL OCCLUSION

IN scenes illuminated with area light sources, regions of partial occlusion or *penumbra* readily occur. Understanding how illumination varies within these regions is both important and difficult. It is important since such an understanding allows us to pick both a suitable sampling strategy, and a good way to compactly represent illumination in the penumbra, using piecewise-polynomial functions for example. These representations can be used for fast high-quality rendering of scenes with area sources, and are important in global illumination calculations (e.g., [16], [21], [34], [12]). However, gaining an understanding of irradiance behavior in the penumbra is difficult because the problem reduces to determining how the visible part of the source changes as one moves from one partially occluded point to another. These changes depend on the interaction of the edges and vertices in the environment; analyzing this geometric interaction is a non-trivial problem. In addition, determining the visible part of the source at any point is expensive if done naively.

The *backprojection* and the *discontinuity mesh* are data structures that permit efficient calculation of the visible portions of a polygonal light source in a penumbra. In polyhedral environments, the irradiance contribution of each portion can be computed analytically using standard techniques. However, the overall irradiance at a point is the sum of all such contributions, and can exhibit visually-

significant variations over a small region. By performing a thorough empirical study of penumbral irradiance behavior, it is possible to glean insights that can be exploited in an efficient approximation. Doing so has allowed us to isolate the causes of multiple extrema in the penumbra. We have also gathered statistics on the frequency of the different configurations affecting penumbral irradiance.

Our study has led to the development of two algorithms that exploit the properties of irradiance in the penumbra. In the first, the number of edges in the discontinuity mesh is reduced without significant deterioration of image quality, while the second algorithm chooses appropriate interpolant degrees (linear or quadratic) again with only moderate quality degradation. Numerical and visual results for both are presented and discussed.

II. PREVIOUS WORK IN SAMPLING AND SHADOW COMPUTATIONS

Approximate and compact representations of illumination or irradiance (impinging light power/area), are useful for the efficient display of illumination for direct lighting and are also necessary for the purposes of global illumination algorithms, such as those developed in radiosity-based approaches (e.g., [6], [17]). In early global illumination algorithms, piecewise constant representations were used for radiosity or irradiance, but it quickly became clear that this representation was insufficient. As an alternative, higher order methods have been since proposed for the solution process (i.e., the light transport phase of global illumination algorithms) with the use of approximation schemes that are of higher degree [31], [16], [34].

A. Observed Properties of Irradiance

Campbell and Fussell [3] observed that irradiance in a penumbral region can exhibit multiple minima and maxima. Numerical optimization was used to determine these critical points. Tampieri [29] and Lischinski et al. [22] segmented the penumbral domain by the mesh generated solely from visual events caused by planar discontinuity surfaces including a source edge or vertex. They then postulated that within each face or cell of this mesh the irradiance varies little. A subsequent adaptive subdivision step was however used when large irradiance discrepancies were observed.

In [9] we proposed that the structure of illumination should be studied in more detail in the hope that a better understanding would lead to more efficient and accurate sampling strategies. Such a structure-driven approach for unoccluded (i.e., shadow-free) environments lit by area light sources was presented in [11]; we conjectured that the illumination from convex polygonal light sources is unimodal, and an effective structured sampling algorithm

George Drettakis is at iMAGIS/GRAVIR-INRIA, BP 53, Grenoble Cedex 9, F-38041, FRANCE. iMAGIS is a joint research project of CNRS, INRIA, UJF, INPG. Part of this work was performed when George Drettakis was a Ph.D. student at the University of Toronto.

Eugene Fiume is at the Department of Computer Science, University of Toronto, Toronto, Ont. CANADA M5S 1A4.

E-mail: George.Drettakis@imag.fr or elf@dgp.toronto.edu




Fig. 1. A Complete Disc. Mesh and Backprojection Instance


based on this conjecture was developed. The algorithm first finds the overall maximum of the irradiance function over the surface, if it exists, and then segments the function into convex and concave regions along two axes passing through the maximum. A mixed quadratic/linear interpolant is fit to the irradiance function satisfying tight and relevant error bounds.

B. Discontinuity Meshing and Backprojections


The visible regions of a polygonal area light source from a point are polygons whose vertices are either formed by the projection of scene edges onto the source, or are vertices of the original light source. A *backprojection instance* at, or induced by, a point P , with respect to a source, is the set of polygons forming the visible parts of the source at that point (e.g., the gray region on the source in Fig. 1). The *backprojection* in a region is a data structure containing the set of ordered lists of emitter vertices and edge pairs such that at every point P in that region, the projection through P of these elements onto the plane of the source form the backprojection instance at P [10].

Given a polygonal light source σ and polygonal scene, the partition of the scene into regions having the same backprojection is the *complete discontinuity mesh* of σ (shown as a yellow mesh in Fig. 1). A region of the complete mesh with the *same* backprojection is a *face* of the mesh. At any point P within a mesh face, the backprojection instance can be efficiently determined by projecting the scene edges in the backprojection structure through the point P to find the coordinates of the relevant points. An example of a scene, its discontinuity mesh, and a backprojection instance (the shaded region on the source) can be seen Fig. 1. The backprojection instance corresponds to the white spot marked P under the drawer.

Early proposals to compute shadows involved numerous techniques dealing with point sources, as well as approximate solutions for linear or area sources (see [33] for a good survey). This research naturally led to the computation of partial discontinuity meshes. In [24] the extremal boundaries were computed, that is the boundary between umbra and penumbra as well as the boundary between penumbra and light for simple geometries. To compute backprojection instances where required, the light source was intersected with the *entire* environment each time, to determine the visible part of the source. The expense of computing ex-



(a)



(b)

Fig. 2. (a) EV and (b) EEE discontinuity surfaces.

act irradiance values in the penumbra was thus prohibitive. Campbell and Fussell [2] first computed shadow boundaries for complex environments using BSP trees from point light sources, and then extended the method to compute the extremal boundaries for area sources in [3]. Chin and Feiner [4] performed a similar computation, and also presented an extension to area sources in [5]. Additional lines of the mesh, interior to the penumbræ, were computed in [22] again using BSP trees, and in [19] using a two-dimensional visibility algorithm. In all these approaches, computing the exact visible portion of the source at a given point involves a visibility computation requiring the intersection of all the scene objects, since the complete mesh has not been computed; thus the backprojection is not unique within each mesh face.

The computation of the complete mesh is performed by casting *discontinuity surfaces* into the environment. These surfaces are of two types: *EV surfaces* which are planar “wedges” caused by the interaction of an edge and a vertex (Fig. 2(a)), and *EEE surfaces* which are ruled quadric surfaces caused by the interaction of three edges (Fig. 2(b)). Algorithms to compute the equivalent problem in computer vision, that of computing the aspect graph, have been proposed among others, by [15] and [14]. An algorithm which computes an equivalent structure for visibility was presented in [30]. An algorithm specifically for shadow computation with good theoretical complexity bounds has been proposed by Stewart and Ghali [27]. A extension and implementation of this approach was presented in [28].

The authors have developed and implemented a fast, practical algorithm for computing the complete discontinuity mesh with backprojections ([8] and [10]). All relevant visual events are properly treated and the algorithm displays fast running times in the number of objects in the

scene, for scenes of moderate complexity. This approach has recently been used to develop a hierarchical global illumination algorithm permitting accurate visibility calculations using backprojections [12].

C. General Discontinuity Meshing and the Irradiance Jacobian

In the work presented in [10] certain special cases (such as EEE surfaces consisting exclusively of edges of the environment) had not been implemented. The work reported here is based on a complete implementation which includes all possible configurations of discontinuity surfaces. We have also developed techniques to treat various degenerate cases of discontinuity surfaces which arise in general environments, permitting the treatment of scenes with arbitrary positioning of the source and scene objects.

Another important addition to structured sampling for unoccluded environments and discontinuity meshing is the use of the irradiance Jacobian as presented by Arvo [1]. This formulation permits the analytic computation of the gradient of irradiance at a cost equivalent to the cost of irradiance $I(p)$ where p is a point of a surface. In particular we can compute and store $\nabla I(p)$. If we wish to determine the derivative dI/dt of irradiance $I(t)$ in a certain direction \bar{u} , we simply perform $\nabla I \cdot \bar{u}$ to obtain the corresponding value. This calculation renders the structured sampling approach of [11] much more efficient and accurate, since the need for numerical approximation of the derivatives is obviated. Arvo presented a formulation for partially occluded points, which allows efficient computation of analytic derivatives, as presented above. In the work presented here the backprojection data structure is used, which provides all the information necessary for the computation in [1].

III. PROPERTIES OF IRRADIANCE FUNCTIONS IN PENUMBRA REGION

Given the complete discontinuity mesh and the backprojection, the exact value of irradiance at any point in the penumbra can be efficiently determined. It thus becomes possible to perform a careful empirical study of irradiance in the penumbra, even for moderately complex scenes. We shall now present the results of an empirical study that isolates key configurations that induce significant variations in penumbral irradiance. We must be aware of these configurations when constructing approximations. Snapshots of such observations are shown in the figures which include the mesh and backprojection geometries, as well as analytically computed irradiance and first derivative values and numerically computed second and third derivatives. Before discussing the results of the experiments, we define some important scene properties.

Backprojection Type. The backprojection of a face is said to be *convex* or *concave* if every instance of the backprojection in that face is itself a convex or concave polygonal subset of the source, respectively. In addition, a backprojection of a face is *disconnected* if every instance of the backprojection in that face consists of more than one polygon, while it is *simple*, if every instance consists of only one




Fig. 3. (a) Convex, (b) concave and (c) disconnected backprojection instances.




Fig. 4. A singular vertex

polygon. An example of a convex backprojection instance is shown in grey on the light source in Fig. 3(a) and a concave example is shown in Fig. 3(b), which are both simple. An example of a disconnected backprojection is shown in Fig. 3(c).

Singularities. As noted in [29] and [22], irradiance along a surface is singular at points at which two surfaces touch. Thus at a vertex in the mesh joining faces of umbra, penumbra and light, the irradiance is multi-valued and is defined as a limit depending on the direction from which the vertex is approached. A singular vertex is shown in Fig. 4, while the graph on the right shows the variation of the irradiance values on three lines joining at the singular vertex.

A. Empirical Characterization of Penumbral Irradiance Behavior

A set of empirical tests were performed within the penumbra for moderately complex scenes. These tests attempted to isolate configurations that cause local extrema. Three influential factors affecting the appearance of extrema were identified:

- *Backprojection type* within the mesh faces of interest; this can cause simple irradiance extrema within a mesh




Fig. 5. (top) Irradiance minimum from disconnected backprojection (bottom) Opposite extrema within a face.

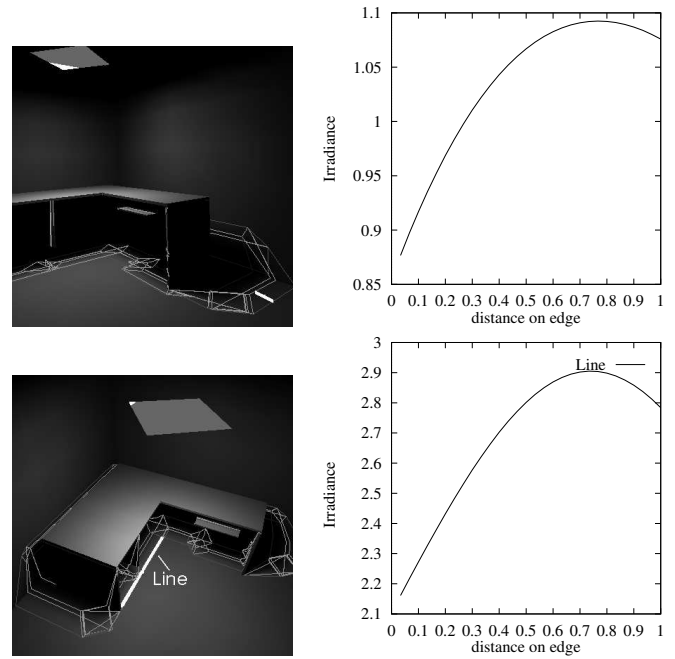


Fig. 6. (top) Interaction with unoccluded regions. (bottom) Constant backprojection instance.

face but also opposite extrema along different directions in the face.

- *Interaction with unoccluded areas*; in particular when unoccluded irradiance increases and the visible part of the source decreases simultaneously or vice-versa.
- *Position of irradiance maximum* along mesh edges with a constant backprojection instance.

In what follows we present examples for each category with an illustration of the backprojection instance shown in grey on the source, and a corresponding graph illustrating the irradiance variation along a line on the floor (shown as a thick white line). We briefly discuss the effects of each configuration on irradiance behavior.

Disconnected or Concave Backprojections. If the backprojection in a face (or along an edge) is not simple, it is likely that there exists a line in the face or an edge for which the irradiance will display one or more maxima. Illumination in a region with a disconnected backprojection containing n polygons is equivalent to illumination from n separate unoccluded light sources, and may have up to n maxima. Non-monotonic behavior in this case may cause “troughs” in the irradiance function, as shown in Fig. 5(top), from the overlap of two “tail” regions [11] of the irradiance function due to a disconnected backprojection with instances containing two polygons. Similar non-monotonic behavior may occur in faces or along mesh edges that have concave backprojections.

Faces with concave or disconnected backprojections, particularly in the presence of light source edges that are very long, can cause the existence of opposite extrema (i.e. a minimum and a maximum) along different directions in a face. An example is shown in Fig. 5(bottom). In this case

the maximum is caused because the corresponding edge in the mesh ends in a light region (see below), while the trough is caused because this face has a concave backprojection.

Interaction with Unoccluded Regions. Consider a line in the mesh that does not lie on the external penumbral boundary between light and shadow but has however one endpoint in light (i.e. an unoccluded region, also called a *light region*). An example is shown as a thick white line in the mesh of Fig. 6(top). The visible area of the source increases along this line, as it goes from penumbra to light, since a smaller portion of the source becomes occluded. As a consequence the irradiance on a line is generally increasing as the region of light is approached. The unoccluded illumination along this line may be increasing or decreasing. If the unoccluded illumination is decreasing (as in Fig. 6(top)), the irradiance in the penumbra, which tends to be an increasing function, will smoothen as it approaches the unoccluded regions. Again, this can result in non-monotonic behavior, although it is necessarily a local maximum.

Position of Maximum from a Constant Backprojection Instance. Consider a scene that is illuminated by a polygonal source, and in which the complete discontinuity mesh of that source has been computed. For some edges in the mesh, the backprojection remains constant along that edge. Thus the analysis used for unoccluded sources in [11] can be used directly. For example, the edge shown as a thick line in Fig. 6(bottom) contains the maximum of the irradiance function from the polygonal source defined by the backprojection instance, shown shaded on the rectangular source. This maximum can be seen in the irradiance plot on the right.




Fig. 7. (a) Image with small source (b) Image with elongated source




Fig. 9. 6 of the 9 light positions used




Fig. 8. Big Scene

B. Scene Statistics and Identification of Non-monotonic Irradiance

A test program has been written which analyses the behavior of irradiance in a set of scenes. The goal is to determine whether irradiance tends to be monotonic within a single face of the discontinuity mesh. This has obvious implications in the necessity for additional adaptive subdivision, but also in the development of error bounds in radiosity calculations [21].

The set of scenes considered consists of a simple desk model in a simple configuration (e.g., Fig. 6) containing 73 polygons and a more complex configuration including drawers (145 polygons) which cause complex visibility interactions (e.g, Fig. 9). For each of the two geometries, the light source was placed in nine different positions. Six of these are shown in Fig. 9. In addition two different light sources were tested, namely a small light source (Fig. 7(a)) and an elongated light source, which behaves in a manner similar to a linear source (Fig. 7(b)) These two source types were tested on both desk geometries (with and without drawers). Finally a larger scene was tested, containing two complex desks and two chairs containing 373 polygons (Fig. 8), which was used to confirm the more exhaustive examples with the multiple source positions. These scenes will also be used to compare the relative performance of the new structured sampling algorithms proposed in later sections (see Section VII-A).

It is important to note that the tests performed using these scenes are of course not definitive. Nonetheless, all types of visibility events occur in these scenes (EV, EEE, D1 [13] and other degenerate configurations) and in addition relatively complex shadow behaviors can be observed (e.g., in the regions below the open drawers or the shadows caused by the back of the chairs). As a consequence we believe that the trends identified in this experimental study are a strong indication of irradiance behavior in the penumbra for interior scenes. We restrict our approach to a single source, since multiple sources present numerous specific issues (as pointed out in [7]).

Given a line embedded in a receiver plane, defined by two endpoints p_t and p_h , we consider the irradiance as it varies along this line as a function $I(t)$ of a single parameter (in a manner similar to that presented previously). In this case we can evaluate $\frac{dI(p_t)}{dt}$ and $\frac{dI(p_h)}{dt}$. If these values are of opposite sign, we consider that the irradiance along the line is *not* monotonic. If the derivative values are of the same sign, we consider the irradiance along this line monotonic.

We have chosen to test the irradiance along all the edges of the mesh on penumbral (i.e., partially occluded) faces, and also along the diagonals of the penumbral faces connecting two vertices of a mesh face not belonging to the same mesh edge. We thus have two tables of statistics for edges (Table I) and for diagonals (Table II), which report *average* values from the 9 different light source positions for “Simple Desk” and “Desk and Drawers”, average of two runs (both geometries) for “Long Source” and “Small Source” and a single run for “Big Scene”. We considered testing faces for monotonicity (i.e. the faces for which either an edge or a diagonal are non-monotonic), but it was observed in test runs that the statistics for faces are very similar to those for the diagonals (Table II).

The monotonicity test is not infallible: edges which are non-monotonic may be ignored because of a change of sign in the derivative in the edge interior. We have however run tests comparing the approach presented above with an exhaustive test of 20 samples of the derivative on each edge, which show that for the scenes in question the simple monotonicity test is accurate for 85% of the tests, which

Scene	e_p	e_{nm}	%NM	f_l	f_d	f_{conc}
Average values over 9 runs						
Simple Desk	721.9	25.1	3.5	23.1	0.9	1.1
Desk Drawers	2386.1	157.6	6.6	33.7	42.2	81.7
Average over 2 runs						
Small Source	906.0	73.5	8.1	28.5	12.5	32.5
Long Source	1449.0	114.5	7.9	27.5	30.5	56.5
Single run						
Big Scene	8661.0	746.0	8.6	78.0	374.0	294.0

TABLE I

MESH EDGE MONOTONICITY OF PARTIALLY OCCLUDED FACES

we judged to be satisfactory. Developing other signatures for nonmonotonicity is an interesting open problem.

In Table I, the field e_p is the average number of edges neighboring at least one mesh face in penumbra, e_{nm} is the number of edges along which the irradiance is non-monotonic, while “%NM” indicates the percentage of edges with non-monotonic irradiance. The breakdown f_l , f_d , f_{conc} , shows the number of the edges with non-monotonic irradiance that neighbor a face respectively in light, with a disconnected backprojection or a concave backprojection. To classify an edge, both neighboring faces are tested, and the edge is designated as either light, disconnected or concave, in that order, if either neighbor is in the appropriate category. For Table II, similar statistics are reported, but for the diagonals *contained* in faces with corresponding properties, with d_p and d_{nm} the total number of penumbral diagonals and those with non-monotonic irradiance respectively, while f_d , f_{conc} , are the number of the diagonals with non-monotonic irradiance inside a face respectively with a disconnected backprojection or a concave backprojection.

From these tables it is clear that irradiance is monotonic along a large majority of edges (more than 92%) of the mesh, consistently, even for different source types and more complex geometry. Similarly, irradiance is monotonic along a large majority of the diagonals. The disparity in Table II between the “Simple Desk” and the “Desk Drawers” scenes is due to the fact that, in the absence of drawers which involves less complex visibility interactions, and consequently larger mesh faces, more significant irradiance variation is observed.

Finally, we can see that a majority of edges or diagonals with non-monotonic irradiance neighbor, or are contained in, faces with disconnected or concave backprojections, as suggested by the empirical study presented above.

C. Discussion

The statistics presented above are meant as a first indication of the behavior of irradiance in penumbral regions. In future work geometric *a priori* determination of which regions of penumbra display non-monotonic behavior should be performed (similar in spirit to the discontinuity “ranking” approach presented recently in [18]).

These first measurements however indicate that irradiance is largely monotonic within regions of equivalent visibility. Such piecewise monotonic functions, especially those

Scene	d_p	d_{nm}	%NM	f_d	f_{conc}
Average values over 9 runs					
Simple Desk	665.3	189.6	28.5	74.4	115.1
Desk Drawers	2181.1	150.0	6.9	43.8	106.2
Average over 2 runs					
Small Source	704.5	61.0	8.7	13.0	48.0
Long Source	1366.5	219.0	16.0	65.0	154.0
Single run					
Big Scene	8116.0	776.0	9.6	438.0	338.0

TABLE II

MESH DIAGONAL MONOTONICITY IN PARTIALLY OCCLUDED FACES

for which the values do not differ significantly, are good candidates for lower order interpolation, specifically linear or quadratic. The experiments presented below will indicate that this is sufficient in general, obviating the need for cubic interpolants as proposed in [26], for a large class of scenes.

In the following section we will present an algorithm for edge elimination, and an algorithm for degree selection. The fact that irradiance in the penumbra is largely monotonic motivates the need to simplify the mesh: if the function is well behaved a coarser subdivision is sufficient. The empirical observations of which factors are important in the penumbra also influenced the construction of the edge elimination algorithm. Similarly the locally well-behaved nature of penumbral irradiance suggests the use of linear interpolation wherever possible, leading to the algorithm for degree selection.

IV. COMBINING UNOCCLUDED STRUCTURED SAMPLING AND DISCONTINUITY MESHING

In [10] and [8] we presented a first attempt at combining the structured sampling approach developed for unoccluded environments and discontinuity meshing. This method consisted of the collection of mesh faces in penumbra into “penumbral groups” which were enclosed in a bounding box and then combined with the subdivision induced by finding the maximum and the inflection points as described in [11]. This approach resulted in reconstruction problems on the boundaries between these bounding boxes of penumbral zones and the unoccluded zones [8].

We have since developed a simpler approach, which appears to avoid these problems. We begin by applying the structured sampling algorithm on the receiver polygon as if it were unoccluded, using the algorithm of [11] (see also Section II-A). This results in a segmentation such as that shown in Fig. 10(a).

In the new approach used here we simply insert the lines of subdivision up to the border of the penumbra and light (Fig. 10(b)). In the quadrilateral regions of the mesh entirely in light we construct bi-quadratic tensor product interpolants. In the regions of penumbra and the regions between regular light regions and penumbra we perform a constrained Delaunay triangulation following [32]. We then construct triangular bi-quadratic interpolants on the resulting triangles (Fig. 10(c)).

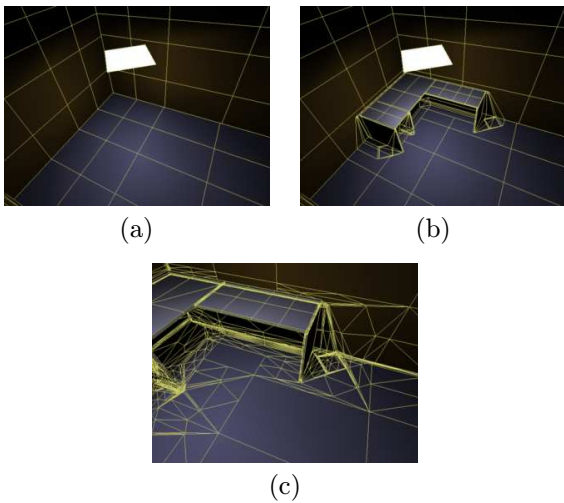


Fig. 10. (a) Unoccluded Structured Sampling (b) Combination with Discontinuity Mesh (c) Triangulation

It is however clear that this construction is too expensive. The size of the discontinuity mesh faces are often small. In addition, in many cases the variation of irradiance in the faces (as well as their size) is so small that linear or even constant approximations are largely sufficient. These facts, supported by the empirical and numerical study presented above, lead us to the development of two algorithms, allowing the reduction of the size of the mesh, and the use of lower degree polynomial approximations presented below.

A. Adaptive Subdivision

Overall, the subdivision effected by the complete discontinuity mesh appears sufficiently fine for the construction of interpolants. In some scenes however, large faces can occur, over which irradiance may vary significantly (e.g., the face containing the thick white edge in Fig. 6(bottom)). As noted by Tampieri [29], adaptive subdivision can be required in such faces.

We apply two simple criteria. The first requires that an edge defined by two points p_t and p_h for which $|I(p_t) - I(p_h)| > \epsilon$, is subdivided. The tolerance ϵ is user-defined. The second requires the subdivision of every edge for which $\dot{I}(p_t)\dot{I}(p_h) = 0$, where $\dot{I}(t)$ is the first derivative. In this manner, edges with non-monotonic irradiance are subdivided. In practice, these criteria result in reasonable triangulations.

Adaptive subdivision is always performed before any simplification or degree selection (see below).

V. INTERPOLANT DOMAIN CONSTRUCTION FOR PENUMBRAL REGIONS

The problem of constructing an interpolant for irradiance in a general scene can be split into the construction of interpolants for regions that are unoccluded, and into the construction of interpolants for regions in partial shadow. The former is treated using the extension of the methods of [11] for unoccluded environments as described above, while the determination of interpolants for penumbral regions is

presented next. There are two aspects to the interpolant definition: the determination of the domains on which the interpolants are defined, and the choice of basis functions, including their degree. We will start with the determination of domains.

A. Constraints on Interpolant Domains

As mentioned above, the complete discontinuity mesh and the accompanying backprojection information is a natural segmentation of the irradiance function over a surface. Edges in this mesh represent discontinuities either of value, or of first or second derivative. Characterizations of these discontinuities can be found in [20], [19], [22], [29].

Value discontinuities occur only where objects touch, and therefore the boundary of such regions delineates areas in which the irradiance function has value zero, because they are completely hidden from the light. Discontinuities of second derivative (or first derivative when degenerate events occur) occur inside regions of penumbra. These discontinuities constitute the majority of edges in the discontinuity mesh. The geometry of the mesh is complicated: it includes highly irregular regions that can be small, concave and with small angles between edges (see Fig. 13(d)). The mesh constrains the construction of interpolant domains, since some of these edges of discontinuity need to be maintained to achieve high quality reconstruction of irradiance.

The relative importance of the discontinuities is difficult to assess without significant numerical computation. Some of the second derivative “jumps” can be small, while others can be quite large. In some cases the effect on the actual irradiance function is more evident (Fig. 11(bottom)) while in others the effect is negligible (Fig. 11(top)). In addition, irradiance in very small regions cannot display extremely large differences in value, because the shape of the polygons in the backprojection instances cannot change much, and neither can the (point-to-area) form-factors that determine the value of irradiance at any point in a face.

To accommodate the highly irregular nature of the faces in the mesh, triangles are selected as the domains over which to construct interpolants. Specifically, a proper triangulation of the penumbral domain is performed. A triangulation $\tau = \{T_0, T_1, \dots, T_n\}$ into n triangles T_i is *proper* if each pair of triangles intersect at a vertex, a complete side or not at all, and the union of all triangles equals the domain (Prenter [25]). In a manner similar to that for irregular regions in light, we use a constrained Delaunay triangulation of the original mesh faces [32]. The basis functions used (presented in Section VI-B) are linear, quadratic or mixed linear/quadratic Lagrange interpolating polynomials.

B. A Mesh Reduction Algorithm

As mentioned above, faces of the discontinuity mesh can be arbitrarily small and may also have edges with high aspect ratios. In addition, the triangulation of concave regions can result in triangles that are very small or that have small angles. For the construction of piecewise smooth interpolant domains, it is desirable to reduce the number

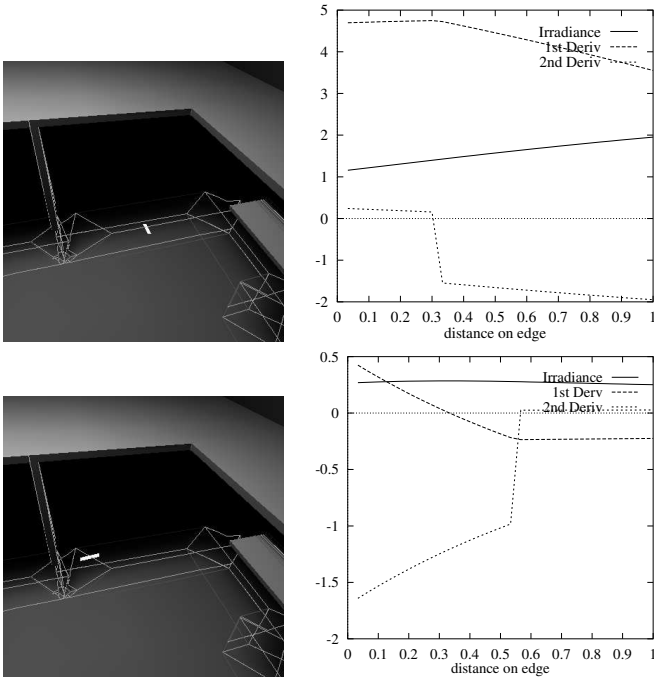


Fig. 11. Discontinuities (top) large “jump” (bottom) small “jump”

of such triangles as much as possible. Larger triangles, and triangles with roughly-equal interior angles are more stable numerically, and in addition provide the benefit of a better theoretical error estimate. Specifically, consider a six-point quadratic interpolant on a triangle, and the irradiance function f . The max norm in an interval $[a, b]$ is defined as follows [25]:

$$\|f(x)\| = \max_{a \leq x \leq b} |f(x)|. \quad (1)$$

Following [25], the error bound with respect to the max norm for an interpolant s_N is:

$$\|f - s_N\| \leq \frac{8 M_3}{\sin \theta} h^3, \quad (2)$$

where h is the longest edge of the triangle, and θ is the smallest angle of the triangle. The constant M_3 is equal to the maximum value of the first, second and third derivative of f within the triangle. In general:

$$M_n = \max\{\|D^1 f\|, \|D^2 f\|, \dots, \|D^n f\|\}. \quad (3)$$

where $D^i f$ is the i 'th derivative of f . In the case of unoccluded illumination, and for some of the faces in the penumbra, the magnitude of the third derivative is many times larger than that of the irradiance itself, rendering Eq. (2) somewhat meaningless (which does not necessarily imply that the approximation is poor). This is shown in Fig. 12(top) where irradiance along the thick white line in the image is plotted in the graph. However, there are other faces in the mesh in which the third derivative is small, and thus the size of the smallest angle θ can play an important role in the quality of the approximation achieved. In Fig. 12(bottom) the absolute value of the third derivative is consistently smaller than the irradiance value.




Fig. 12. Absolute values of third derivative in the mesh faces along the thick white lines.

B.1 An Area-Based Edge Elimination Algorithm

For the reasons outlined above it is desirable to eliminate all overly small faces in the mesh. A smaller mesh size is desirable in general, among other reasons because it allows the irradiance function to be represented with a smaller number of triangular interpolants. This allows faster rendering and is also important when such a representation is used for light transfer simulation (see [12], [8]). In the results it will be shown that a large mesh reduction with negligible error is achievable.

Our guiding strategy is to remove all faces that have area less than `area-tol`% of the largest face, where `area-tol` is a user defined tolerance. When eliminating edges from the mesh, corresponding faces are deleted. It is therefore necessary to maintain the backprojection information, together with the geometry of the deleted faces in the resulting merged face. This information can be discarded after the construction of the interpolants if the exact solution is no longer desired.

Eliminating Concavities. It is often the case that the tip of a smaller face will bite into a larger face, creating a concavity. For the reasons outlined above, it is desirable to eliminate such faces. Each concave face is visited, and the vertices for which the two edges on the face form an angle greater than 180 degrees are identified. Edges are removed from the concave face if the corresponding small face has area below the tolerance. If all the faces around the vertex can be removed, the remaining unconnected edge is also removed from the mesh.

Eliminating Small Faces. After the concave faces have been treated, a number of small faces that have area less than the maximum may still remain. For each such face




Fig. 13. (a) Image and (b) mesh before reduction (c) detail image (d) detail mesh before reduction

the following procedure is applied:

```

remove_small_face( Face f )
{
    sort all edges e of f
      by area of neighboring face
    for each edge e in sorted list do
      if can_remove( e ) then
        remove_edge( e );
        return;
      endif
    done
}

```

The procedure `can_remove(e)` determines whether the removal of the edge is possible. The following rules are applied:

1. No edge of the boundary between umbra and penumbra is removed.
2. No edge of the boundary between umbra and light is removed.
3. If the removal of an edge results in a convex face becoming concave, it will not be removed.

After this procedure is applied, all remaining unconnected edges are removed from the remaining face. After an edge is deleted, the tail and head vertex are searched to determine if the edge extended into the neighboring faces. If it did, the neighboring edges are removed, if such a removal does not violate rules 1-3. The removal process (concavity and small face removal) is repeated until all faces that are smaller than the tolerance are removed, or until no more faces can be removed.




Fig. 14. (a) Image and (b) mesh after 56.0% reduction (c) detail image (d) detail of reduced mesh

Polygon rendering hardware on an SGI Indigo 2 XZ was used to generate images in Fig. 13 to 15. In Fig. 13 part of “Big Scene” is shown using the combined structured and discontinuity mesh algorithm (Section IV) before mesh reduction, while in Fig. 14, the same scene and mesh are shown after reduction. Although the mesh has been reduced by 56.0% of the original faces before reduction, the quality of interpolation is still good. The reduction of the mesh is shown in detail in Fig. 14(d), compared to Fig. 13(d). The images with a reduced mesh, even when looking at details (Fig. 14(c) vs. Fig. 13(c)) are still of acceptable quality.

VI. CONSTRUCTING INTERPOLANTS FOR PENUMBRAL REGIONS

Once the triangular domains have been constructed, we have to choose suitable basis functions and then calculate the coefficients to construct the interpolant. The empirical and numerical analysis presented in Section III suggests that for many cases in the penumbra, linear interpolation of irradiance is sufficient. Thus the interpolants constructed by our algorithm are chosen to be of low degree: linear, quadratic or mixed.

Because the complete mesh has been computed, the umbral regions are well defined. On these domains, constant basis functions are used with a value of zero. For this step, the computation of the complete mesh is a necessity, since boundaries of the umbra are often (EEE-induced) curved edges, which were not computed by previous discontinuity meshing algorithms (e.g., [22], [19]) although they have been treated for a different application in [30].

To achieve the construction of mixed linear/quadratic

interpolants, it is first necessary to characterize the edges of the mesh and the edges interior to the faces with a required degree of interpolation. Each edge or face is designated as linear or quadratic, and the appropriate basis function is assigned to each triangle of the face, maintaining C^0 continuity.

A. Selecting a Degree for the Triangular Domains

The choice of degree required on an edge of the discontinuity mesh (DM) or a triangle edge in the interior of a DM face is determined by whether or not irradiance along it is monotonic, and whether the difference between irradiance values at the endpoints is larger than a user-defined tolerance `linear-tol`. The algorithm first determines monotonicity on each edge using the same approach as that presented in Section III-B. An edge is then marked “constant” if the difference of irradiance at its endpoints is zero and “linear” if it is less than `linear-tol`. Otherwise it is marked “quadratic”.

As can be seen from the statistics presented in Section III-B, a large proportion of non-monotonic edges are adjacent to faces with disconnected and concave backprojections. Edges neighboring a light face (i.e. a face of the mesh in an unoccluded region), and mesh edges that include the maximum of irradiance can also be non-monotonic. For edges of triangles interior to faces, similar properties hold. Specifically, faces with concave or disconnected backprojections, as well as faces neighboring light regions often display non-monotonic behavior.

Once each edge of the triangular domains has been classified as “constant”, “linear” or “quadratic”, a triangular basis function is selected and the coefficients (irradiance values) are calculated at the triangle vertices and appropriate interior points. In Fig. 15 the same scene as that of Fig. 13 is rendered after the application of the degree selection algorithm. Notice that the visual quality of interpolation is still high despite the fact that 58.3% of the interpolants are linear. The red lines show the edges of the mesh that have been assigned linear interpolants, while the green lines show edges with quadratic interpolants. The choice of basis functions and their construction is described in the following section.

B. Basis Function Design

The basis functions chosen to interpolate the irradiance are triangular Lagrange polynomials. For triangles in which all three edges are linearly interpolated, the method of plates is used to construct a linear basis function over the triangle. For this interpolant the formal error bound is given by [25]:

$$\max |f(p) - s_N(p)| \leq 4M_2h^2 \quad (4)$$

where f is the irradiance function and s_N is the interpolant, and h and M_2 are as in Eqs. (2) and (3). For triangles on which all edges are of degree 2, a six-point interpolant is used [25]. The error bound for this interpolant was given in Eq. (1).




Fig. 15. Image and mesh degree selection (58.3% of edges are linear shown in red)




Fig. 16. Linear/quadratic basis functions.

As a consequence of degree selection, there are triangles for which some edges are linear and some are quadratic (see Fig. 16(a) which follows the model of [25]). A special linear/quadratic basis function was designed to guarantee value continuity across such interpolants. For a given vertex, the six-point bi-quadratic interpolant is used, which interpolates the value 1 at the vertex of interest and the line passing through the values 0, $\frac{1}{2}$, 1 along the linear edge. This basis function is depicted in Fig. 16(b). These samples are of course reused on neighboring triangles.

For triangles with one linear edge, two linear/quadratic basis functions are used at the vertices of the linear edge, while the remaining vertex is assigned a quadratic basis function, and two quadratic basis functions along the quadratic edges. A total of 5 function (irradiance) values are required for the triangle with one linear edge, since we require one value at each vertex, plus two values at the midpoints of the edges characterized as quadratic.

For triangles with two linear edges, all vertices have linear/quadratic basis functions, and an additional quadratic basis function is defined at the point along the quadratic edge. A total of 4 function values are required for this interpolant.

C. An Interpolant for Triangles at a Singularity

A typical singular vertex will have many edges joining at that point, as described in Section III. The value at the vertex is defined as a limit of the function as it approaches the singular point, and therefore it is multi-valued. To represent this using interpolants, as suggested by Tampieri [29], we use degenerate tensor product interpolants. Nonetheless, the multiple values of irradiance at this vertex must be approximated. To achieve this, the total angle θ_{tot} of the triangles neighboring the vertex is first computed.




Fig. 17. Singular vertex interpolant construction

The value of the unoccluded illumination at the singular vertex L_{max} is calculated next. For triangle i , whose edges joining at the singularity form angle θ_i , (see Fig. 17) the values of irradiance for the bi-quadratic tensor products along the edge corresponding to the singular point are assumed to vary in the interval:

$$\left[L_{max} \sum_{i=0}^{j-1} \frac{\theta_i}{\theta_{tot}}, L_{max} \sum_{i=0}^j \frac{\theta_i}{\theta_{tot}} \right]. \quad (5)$$

This is an approximation, since it would be necessary to compute the limit values of irradiance along each edge to determine the exact range of irradiance in each triangle. The midpoint vertex in the tensor product is assigned the average of these two values. In this fashion a continuous representation is constructed for the irradiance around a singularity.

A special case occurs when the light source touches a different object. In this case we simply displace the vertex slightly so that we can apply the analytical point-to-area form-factor (e.g., [1]) to determine L_{max} .

D. Error Bounds for the Triangular Interpolants

The theoretical error bounds for the linear and quadratic triangular interpolants have been given in Eqs. (2) and (4). These bounds are conservative, since they depend on the maximal magnitude of the derivatives within the domain of interest. In some faces of the discontinuity mesh the derivatives are sufficiently small to permit these bounds to be meaningful. Identifying these cases however is expensive, and not a practical way of assessing quality.

For faces with monotonic irradiance, a simple error bound is given from the maximum difference between irradiance values at the three points of the polygon. The error bound B of the interpolant s_N , on triangles over which irradiance is monotonic, is given as:

$$B = \max |f(p_i) - f(p_j)|, \quad i, j = 1, 2, 3. \quad (6)$$

It is important to note that since we do not guarantee correct monotonicity characterization, in terms of our algorithm B is a heuristic rather than a strict bound.

Error bounds have not been strictly established for faces with non-monotonic irradiance for which adaptive subdivision has not been performed. However, for such faces with

a single extremum, the quality of the interpolant can be estimated closely by the maximum of B in Eq. (6) and the maximum difference of the irradiance values of the interior points along each edge of the triangle used in the quadratic interpolant construction.

VII. NUMERICAL TESTS AND QUALITY EVALUATION

To evaluate the quality of the algorithms for mesh reduction and degree selection we ran the algorithms on the “Simple Desk” and “Desk Drawers” scenes for the 9 light positions as described in Section III-B, as well as the two geometries for “Small” and “Long” light source configurations and finally “Big Scene” with two desks and two chairs. All images used for the numerical evaluation (Tables III to VII) were computed using ray-casting, with the value of the visible point at each pixel being determined by evaluating the interpolant of calculating the exact backprojection (for the reference images). To measure error we use two error metrics: an area-weighted object-space square root error and an image-based square root error.

For the object-space error we compute a set of sampling points $\{p_j\}$ on each face f_i of the mesh, both in penumbra and in light. Define $I(p)$ to be the irradiance on a surface at point p and $\hat{I}(p)$ the approximation constructed by piecewise linear/quadratic interpolation. We have F faces in the scene and each face f_i has area A_i . The object space error ϵ_{os} is thus given as:

$$\epsilon_{os} = \sqrt{\frac{1}{\sum A_i} \sum_i \frac{A_i}{n} \sum_{j=0}^n (\hat{I}(p_j) - I(p_j))^2} \quad (7)$$

The error metric ϵ_{os} has the same units as irradiance (power per area). We also compute a reference image with analytic radiance values $E(i, j)$ and an approximate image $\hat{E}(i, j)$ using the interpolants. The image-based error ϵ_{is} for a $n_i \times n_j$ image, is given as:

$$\epsilon_{is} = \sqrt{\frac{\sum_i^{n_i} \sum_j^{n_j} (E(i, j) - \hat{E}(i, j))^2}{n_i n_j}} \quad (8)$$

The error metric ϵ_{is} is in pixel value differences and thus varies between 0 and 255. We compute ϵ_{is} for three difference viewpoints, the first is the view shown in Fig. 9, the second is similar to the image in Fig. 13 and a third which is a view on the other side of the desk. For “Big Scene” the view of Fig. 8 is also used.

A. Results for Mesh Reduction and Degree Selection Algorithms

Results are presented next in Tables III to VII. Table III gives the results for the simple combination of structured sampling with discontinuity meshing, Tables IV,V show the statistics for the mesh reduction algorithm, while Tables VI, VII shows the results for the degree selection approach.

In Tables IV,V, n_f is the original number of faces in the mesh, and n'_f is the reduced number, ϵ_{os} and ϵ_{is} are as defined above and % *Red.* is the percentage reduction of

Scene	ϵ_{os}	ϵ_{is}
Average over 9 runs		
Simple Desk	0.019	0.461
Desk Drawers	0.025	0.664
Average over 2 runs		
Small Source	0.015	0.716
Long Source	0.024	1.623
Single run		
Big Scene	0.004	0.851

TABLE III
ERROR FOR COMBINED ALGORITHM

Scene	n_f	n'_f	%Red.	ϵ_{os}	ϵ_{is}
Average over 9 runs					
Simple Desk	344.9	244.1	29.2	0.016	0.052
Desk Drawers	1025.4	673.0	34.4	0.019	0.076
Average over 2 runs					
Small Source	421.5	260.5	38.2	0.018	0.444
Long Source	647.0	385.0	40.5	0.106	0.808
Single run					
Big Scene	3840.0	2529.0	34.1	0.016	1.260

TABLE IV
MESH REDUCTION: AREA TOLERANCE 0.01 (1%)

Scene	n_f	n'_f	%Red.	ϵ_{os}	ϵ_{is}
Average over 9 runs					
Simple Desk	344.9	159.9	53.6	0.021	0.067
Desk Drawers	1025.4	375.0	63.4	0.021	0.102
Average over 2 runs					
Small Source	421.5	187.0	55.6	0.028	0.854
Long Source	647.0	237.0	63.4	0.099	1.129
Single run					
Big Scene	3840.0	1689.0	56.0	0.022	3.946

TABLE V
MESH REDUCTION: AREA TOLERANCE 0.09 (9%)

the number of faces. In Tables VI,VII, n_{eq} is the number of edges with quadratic interpolants, n_{el} is the number of edges with linear interpolants, and % *Lin.* the percentage of edges with linear interpolants.

The mesh reduction achieved is satisfactory (between 29% to 63%), while the error is globally low. Similarly, the number of edges characterized as linear is high (from 27% to 69.1%), allowing the use of cheaper, lower degree polynomial interpolants. Object-space error ϵ_{os} is low, growing slightly more than in other cases for mesh reduction in the case of the long light source (Table IV), for which the tolerance value results in higher mesh reduction. Image space error ϵ_{is} is also low, since it is less than a unit RGB pixel value for almost cases of mesh reduction or degree selection for the simple scene, and in the order less than 4 RGB pixel values for mesh reduction of the big scene (Table V).

We next present a first visual comparison of the images presented in Sections VI and V by showing the difference images multiplied by a factor of 20. As can be seen in Figs. 18 and 19 the differences are limited and small in magnitude (a dark pixel signifies no difference, and a to-

Scene	n_{eq}	n_{el}	% Lin.	ϵ_{os}	ϵ_{is}
Average over 9 runs					
Simple Desk	1049.7	396.0	27.4	0.002	0.051
Desk Drawers	1909.2	1526.4	44.4	0.004	0.075
Average over 2 runs					
Small Source	863.0	993.5	53.5	0.012	1.348
Long Source	1426.5	866.0	37.8	0.012	0.813
Single run					
Big Scene	4736.0	6630.0	58.3	0.004	0.865

TABLE VI
DEGREE SELECTION TOLERANCE 0.001

Scene	n_{eq}	n_{el}	% Lin.	ϵ_{os}	ϵ_{is}
Average over 9 runs					
Simple Desk	1780.2	1111.1	38.4	0.019	0.063
Desk Drawers	2982.0	3889.3	56.6	0.024	0.097
Average over 2 runs					
Small Source	1702.5	2010.5	54.1	0.030	1.348
Long Source	2327.5	2257.5	49.2	0.121	1.124
Single run					
Big Scene	2443.0	8923.0	89.3	0.021	0.955

TABLE VII
DEGREE SELECTION TOLERANCE 0.008

Scene	t_{tot}	t_m	t_{bp}	t_{ss}	$t_r^{0.01}$	$t_r^{0.09}$
Average over 9 runs						
Simple Desk	21.40	15.43	2.98	0.79	0.16	0.16
Desk Drawers	116.56	73.83	33.68	1.20	1.04	0.74
Average over 2 runs						
Small Source	23.55	18.61	1.88	0.97	0.23	0.17
Long Source	56.56	40.46	10.75	0.80	0.49	0.36
Single run						
Big Scene	649.22	440.60	175.05	2.14	13.33	9.24

TABLE VIII
TIMING RESULTS

tally white pixel a difference of 255 pixel levels in all three channels).

Finally, timing results are presented in Table VIII. All timings are in CPU seconds on an SGI Indy R5000 processor running at 150Mhz. The value t_{tot} is the total time spent to create the mesh, triangulate and construct the interpolants, t_m is the mesh construction time, t_{bp} is the backprojection calculation time, t_{ss} is the cost of structured sampling (as in [11] but using irradiance gradients [1]), $t_r^{0.01}$ is the cost of the edge removal algorithm for an area tolerance of 0.01 (1%) and $t_r^{0.09}$ is the removal time for area tolerance of 0.09 (9%). As can be seen from these statistics, the cost of the edge removal algorithm is negligible compared to the total cost of the algorithm.

B. Discussion of Numerical and Visual Results

The mesh reduction algorithm has presented good results for the scenes tested. For satisfactory mesh reduction (30-60%), the increase in error is in general minimal, indicating that the complete discontinuity mesh is much larger than required for satisfactory reconstruction. However, there are cases in which the reconstruction quality can




Fig. 18. Difference images (x20) for: (a) simple interpolant (b) reduced mesh.




Fig. 19. Difference image (x20) after degree selection.

degrade, particularly when small faces are left in the mesh which force the creation of small or elongated triangles. Further geometric manipulation of the mesh can be used to eliminate these artifacts. One issue that is more difficult to address is that of animation. No provision is currently made for consistency in mesh reduction, and thus over animated sequences flickering can occur as the mesh changes from frame to frame. This should be addressed in the context of a more general incremental meshing algorithm for animation.

For degree selection, the results are also encouraging. With a good percentage of linear interpolants (30-70%) the increase in error is small, both in object and image space. Nonetheless, the method can be improved by incorporating some criterion based on the possible visible impact of the degree reduction, to avoid occasionally objectionable artifacts. This will require the use of perceptual error metrics.

VIII. SUMMARY AND DISCUSSION

An empirical and numerical study of irradiance in penumbral regions has been presented. Such a study was previously impractical due to the expense of computing irradiance values in the penumbræ. The use of the discontinuity mesh and the backprojection now makes such a study possible. It was found that in the majority of cases, irradiance in the penumbra is monotonic and thus amenable to reconstruction by linear or quadratic interpolants. Configurations that cause the appearance of extrema or irregular behavior were characterized. This study offers a better understanding of irradiance behavior in regions of partial occlusion, and guided the construction of an interpolating,

interpolating, piecewise polynomial representation.

For the construction of the interpolant domains, the complete discontinuity mesh is used as a starting point. The faces of the mesh are triangulated, and the irradiance information is stored compactly with the mesh. The observations made from the empirical study suggested that many of the edges in the mesh are not actually required for satisfactory reconstruction of illumination in the penumbra. A mesh reduction algorithm is thus introduced, based on the removal of faces with small area and faces that cause concave regions in the mesh.

The observations of the empirical study also suggest that in many cases linear interpolation is sufficient for illumination reconstruction in the penumbra. An algorithm was presented which characterizes the edges in the mesh and the triangulation as requiring linear or quadratic degree polynomials to achieve high quality reconstruction. A set of basis functions was designed that allows the use of mixed degree polynomials for reconstruction.

Numerical tests were performed for a suite of moderately complex environments in which complicated shadow structures appear. The results show that both the mesh reduction and the degree selection algorithm can be applied without significant degradation of quality in the reconstruction.


This paper is a first attempt at comprehending irradiance behavior in penumbral regions. Much more work remains to be done. More detailed studies are needed of the geometric conditions leading to irradiance extrema in the penumbra. This will hopefully lead to a priori algorithms which will allow the casting only of those discontinuity surfaces which contribute to significant illumination changes. In this manner the mesh will be simplified overall, and the meshing will be computationally cheaper. Such work will also result in much more reliable and effective error bounds, which are indispensable for global illumination algorithms [21]. The hierarchical global illumination algorithm incorporating discontinuity meshing and backprojections presented in [12] can use the mesh reduction approach in a straightforward manner. A first attempt at simplification in the presence of multiple sources was presented in [7], and evidently more work needs to be performed in combining the different simplification strategies, in particular for the application of discontinuity meshing to global illumination [23], [12].

IX. ACKNOWLEDGEMENTS


The authors would like to thank the various members of the University of Toronto Dynamic Graphics Project for their contributions to the discontinuity meshing system, notably Rob Lansdale, James Stewart and Pierre Poulin. The referees are thanked for the careful review of the manuscript and the many suggestions and corrections which greatly improved the quality of this paper. The research in this paper was funded in part by the University of Toronto, ITRC (Ontario), NSERC (Canada) and ERCIM (European Commission).

REFERENCES

- [1] James Arvo. The irradiance jacobian for partially occluded polyhedral sources. In *Computer Graphics Proceedings, Annual Conference Series: SIGGRAPH '94* (Orlando, FL). ACM SIGGRAPH, New York, July 1994.
- [2] A. T. Campbell, III and Donald S. Fussell. Adaptive mesh generation for global diffuse illumination. In Forest Baskett, editor, *Computer Graphics (SIGGRAPH '90 Proceedings)*, volume 24, pages 155–164, August 1990.
- [3] A. T. Campbell, III and Donald S. Fussell. Analytic illumination with polygonal light sources. Technical Report R-91-15, Dept. of Computer Sciences, Univ. of Texas at Austin, April 1991.
- [4] Norman Chin and Steven Feiner. Near real-time shadow generation using BSP trees. In Jeffrey Lane, editor, *Computer Graphics (SIGGRAPH '89 Proceedings)*, volume 23, pages 99–106, July 1989.
- [5] Norman Chin and Steven Feiner. Fast object-precision shadow generation for areal light sources using BSP trees. In David Zeltzer, editor, *Computer Graphics (1992 Symposium on Interactive 3D Graphics)*, volume 25, pages 21–30, March 1992.
- [6] Michael F. Cohen and Donald P. Greenberg. The hemi-cube : A radiosity solution for complex environments. *Computer Graphics*, 19(3):31–40, July 1985. Proceedings SIGGRAPH '85 in San Francisco (USA).
- [7] George Drettakis. Simplifying the representation of radiance from multiple emitters. In *Proceedings of 5th EG Workshop on Rendering*, Darmstadt, Germany, June 1994.
- [8] George Drettakis. *Structured Sampling and Reconstruction of Illumination for Image Synthesis*. PhD thesis, Department of Computer Science, University of Toronto, January 1994. PhD Thesis, CSRI Tech. Report No. 293:ftp:ftp.csri.toronto.edu:csri-technical-reports/293.
- [9] George Drettakis and Eugene Fiume. Concrete computation of global illumination using structured sampling. *Third Eurographics Workshop on Rendering*, May 1992.
- [10] George Drettakis and Eugene Fiume. A fast shadow algorithm for area light sources using back projection. In *Computer Graphics Proceedings, Annual Conference Series: SIGGRAPH '94* (Orlando, FL), pages 223–230. ACM SIGGRAPH, New York, July 1994.
- [11] George Drettakis and Eugene L. Fiume. Accurate and consistent reconstruction of illumination functions using structured sampling. *Computer Graphics Forum (Proc. of Eurographics '93)*, 13(3):273–284, September 1993.
- [12] George Drettakis and Francois Sillion. Accurate visibility and meshing calculations for hierarchical radiosity. In *Proceedings of Seventh Eurographics Workshop on Rendering in Porto, Portugal*. Eurographics, June 1996.
- [13] Sherif Ghali and A. James Stewart. A complete treatment of d1 discontinuities in a discontinuity mesh. *Proceedings of Graphics Interface '96*, pages 122–131, May 1996.
- [14] Ziv Gigus, John Canny, and Raimund Seidel. Efficiently computing and representing aspect graphs of polyhedral objects. *IEEE Trans. on Pat. Matching & Mach. Intelligence*, 13(6), June 1991.
- [15] Ziv Gigus and Jitendra Malik. Computing the aspect graph for the line drawings of polyhedral objects. *IEEE Trans. on Pat. Matching & Mach. Intelligence*, 12(2), February 1990.
- [16] Steven Gortler, Michael F. Cohen, and Philipp Slusallek. Radiosity and relaxation methods. *submitted for publication*, 1993.
- [17] Pat Hanrahan, David Saltzman, and Larry Aupperle. A rapid hierarchical radiosity algorithm. *Computer Graphics*, 25(4):197–206, August 1991. SIGGRAPH '91 Las Vegas.
- [18] Stephen Hardt and Seth Teller. High-fidelity radiosity rendering at interactive rates. In *Proceedings of Seventh Eurographics Workshop on Rendering in Porto, Portugal*. Eurographics, June 1996.
- [19] Paul Heckbert. Discontinuity meshing for radiosity. *Third Eurographics Workshop on Rendering*, pages 203–226, May 1992.
- [20] Paul S. Heckbert and James M. Winget. Finite element methods for global illumination. Technical Report UCB/CSD 91/643, Computer Science Division (EECS), University of California, July 1991.
- [21] Dani Lischinski, Brian Smits, and Donald P. Greenberg. Bounds and error estimates for radiosity. In *Computer Graphics Proceedings, Annual Conference Series: SIGGRAPH '94* (Orlando, FL), pages 67–74. ACM, July 1994.
- [22] Dani Lischinski, Filippo Tampieri, and Donald P. Greenberg. Discontinuity meshing for accurate radiosity. *IEEE Computer Graphics and Applications*, 12(6):25–39, November 1992.
- [23] Dani Lischinski, Filippo Tampieri, and Donald P. Greenberg. Combining hierarchical radiosity and discontinuity meshing. In *Computer Graphics Proceedings, Annual Conference Series: SIGGRAPH '93* (Anaheim, CA, USA), pages 199–208. ACM, August 1993.
- [24] Tomoyuki Nishita and Eihachiro Nakamae. Continuous tone representation of three-dimensional objects taking account of shadows and interreflection. In B. A. Barsky, editor, *Computer Graphics (SIGGRAPH '85 Proceedings)*, volume 19, pages 23–30, July 1985.
- [25] P. M. Prenter. *Splines and Variational Methods*. John Wiley & Sons Inc, New York, 1989.
- [26] David Salesin, Dani Lischinski, and Tony DeRose. Reconstructing illumination functions with selected discontinuities. *Third Eurographics Workshop on Rendering*, pages 99–112, May 1992.
- [27] A. James Stewart and Sherif Ghali. An output sensitive algorithm for the computation of shadow boundaries. In *Canadian Conference on Computational Geometry*, pages 291–296, August 1993.
- [28] A. James Stewart and Sherif Ghali. Fast computation of shadow boundaries using spatial coherence and backprojections. In Andrew Glassner, editor, *Proceedings of SIGGRAPH '94 (Orlando, Florida, July 24–29, 1994)*, Computer Graphics Proceedings, Annual Conference Series, pages 231–238. ACM SIGGRAPH, ACM Press, July 1994. ISBN 0-89791-667-0.
- [29] Filippo Tampieri. *Discontinuity Meshing for Radiosity Image Synthesis*. PhD thesis, Department of Computer Science, Cornell University, Ithaca, New York, 1993. PhD Thesis.
- [30] Seth J. Teller. Computing the antipenumbra of an area light. *Computer Graphics*, 26(4):139–148, July 1992. Proceedings of SIGGRAPH '92 in Chicago (USA).
- [31] Roy Troutman and Nelson L. Max. Radiosity algorithms using higher order finite element methods. In *Computer Graphics Proceedings, Annual Conference Series: SIGGRAPH '93* (Anaheim, CA, USA), pages 209–212. ACM SIGGRAPH, New York, August 1993.
- [32] Marc Vigo. An incremental algorithm to construct restricted delaunay triangulations. Technical Report LSI-95-43-R, Departament de Llenguatges i Sistemes Informatics, Universitat Politècnica de Catalunya, Barcelona, Spain, 1995.
- [33] Andrew Woo, Pierre Poulin, and Alain Fournier. A survey of shadow algorithms. *IEEE Computer Graphics and Applications*, 10(6):13–32, November 1990.
- [34] Harold R. Zatz. Galerkin radiosity: A higher-order solution method for global illumination. In *Computer Graphics Proceedings, Annual Conference Series: SIGGRAPH '93* (Anaheim, CA, USA), pages 213–220. ACM SIGGRAPH, New York, August 1993.



George Drettakis received his Ptychion (B.Sc. equivalent) in Computer Science at the University of Crete, Greece (1988). He received an M.Sc. and Ph.D. in Computer Science at the University of Toronto, Canada, in 1990 and 1994 respectively. He received an ERCIM post-doctoral fellowship which was performed at INRIA, (France), UPC, (Spain) and GMD (Germany) in 1994-1995. He holds the position of researcher at INRIA Rhône-Alpes, Grenoble, where he works in the iMAGIS/GRAVIR-INRIA project since October 1995. His research interests include global illumination, visibility and shadow calculations, sampling, error estimation and the treatment of complex environments. He is a member of ACM and Eurographics.



Eugene Fiume is in the Department of Computer Science at the University of Toronto, where he has been on faculty since 1987. Over that period he has also been an NSERC University Research Fellow, and co-directs the Dynamic Graphics Project. Prior to his faculty appointment, he was Maitre Assistant at the University of Geneva, Switzerland. He holds Ph.D. and M.Sc. degrees from the University of Toronto,

and a B.Math. from the University of Waterloo. He was a Visiting Professor at the University of Grenoble, France in 1996. He is currently Associate Director of CSRI the University of Toronto, and is a member of the Scientific Advisory Board of GMD, Germany. He is also a consulting Senior Scientist at Alias|Wavefront Inc. His research interests include most aspects of realistic computer graphics, including computer animation, modelling natural phenomena, illumination, image processing, sampling and filtering, and graphics software systems. He is a member of the ACM and Eurographics, and is an associate member of the IEEE.